

**SC15  
NASA Research Exhibit  
Booth 333**

**Abstracts are for presentations to be given at the SC15 NASA Exhibit,  
At SC15 in Austin, TX  
November 16-19, 2015**

**Abstracts will also be posted to NASA@SC15 website  
beginning 11/16/15**

## Off-Design Optimization of Flexible Wings with Distributed Flaps

### Overview

An aircraft's weight varies greatly as fuel burns throughout cruise putting the wing "off-design" for the majority of the flight. This situation is exacerbated by newer, more flexible wings which unbend significantly during cruise as fuel burns and the load decreases. This work examines the potential of a full-span distributed flap system to adaptively re-shape the wing in order to recover the lost wing performance and keep the wing operating near peak efficiency throughout the aircraft's mission.

### Project Details

This work involves a nested, adjoint-driven optimization process combining aerodynamic shape optimization with structural modeling. It optimizes deflections of an array of 48 distributed flaps to reduce lift induced drag at multiple operating conditions in the mission. Since the flaps alter the lift distribution over the wing, the method deforms the flexible wing in each candidate design in response to this varying load. The aero-static computation of wing shape is done through an iterative aero-structural computation for the bending and twist of the wing. The structural analysis includes varying fuel loads, structural and engine weights, engine thrust and other important factors in the analysis.

### Results and Impact

This work developed a fast and powerful methodology for aero-structural shape optimization for highly flexible wings which is many times more computationally efficient than traditional approaches and has the benefit of full automation. This approach was employed to examine the potential benefit of actively managing the load distribution of highly flexible wings through adaptive deflection of a full-span flap system.

- The full-span flap system was able to manage wing's lift distribution to achieve near optimal performance throughout the entire mission profile.
- The flap system also showed great promise for reducing penalties from the strong wing shock when the aircraft was operating at off-design conditions.

### Role of High-End Computing (Why HPC Matters)

While the new approach is roughly 5 times more computationally efficient than typical coupled aero-structural analysis, high-fidelity aerodynamic shape optimization is very compute intensive. A typical design may examine dozens to hundreds of potential candidates before finding the optimal wing, severely taxing most large computers. The extreme computational power available in NASA's Advanced Supercomputing division made it possible not only to perform these design exercises, but also to quantify the benefit through comparison with baseline configurations, conventional stiff-wing aircraft, best-possible reference configurations and other key reference points. In all about 2 dozen full aero-structural optimizations were performed equating to thousands of individual CFD simulations.

### What's Next

The ability of the distributed flap system to actively alleviate wave drag was an unexpected discovery in this work. We are actively investigating the degree to which the flap system can reduce this drag penalty in so-called "dash cruise" cases where the aircraft is purposely flown faster than its design Mach number.

Presented at SC15 by:

**PI** Michael J. Aftosmis, NASA Ames Research Center, michael.aftosmis@nasa.gov

Co-Principal Investigator:

David L. Rodriguez, NASA Ames Research Center, david.l.rodriguez@nasa.gov

Exhibit Theme: Aeronautics

## **Ensuring Safe Passage of the Space Launch System Through the Speed of Sound**

### **Overview**

During the ascent flight phase of the Space Launch System (SLS), NASA's next heavy-lift launch rocket, the vehicle may encounter flow-induced oscillatory forces (buffet) which can excite the vehicle structure. If these forces exceed predicted design loads, this buffeting could result in a catastrophic end to flight. Experimental data is typically the means used to develop loading models to evaluate structural responses to flow field dynamics, during flight in the near speed of sound (transonic) range. This project used computational fluid dynamic (CFD) simulations of transonic flight conditions to predict the time-varying surface pressures and resulting oscillatory forces. Comparisons to experimental wind tunnel data showed good agreement with peak RMS location, with dominant frequency at that location predicted within 20%.

### **Project Details**

A Delayed Detached Eddy Simulation (DDES) model was used to computationally obtain the unsteady flow about the entire SLS vehicle. Surface pressure time histories were monitored at over 600 locations on the SLS and compared to experimental data from both the NASA Langley Transonic Dynamics Tunnel (TDT) and the NASA Ames Unitary Plan Wind Tunnel (AUPWT).

Sensitivity studies of computational mesh density and time step size were performed to quantify the accuracy of computed data. The root-mean-square (RMS) of pressure coefficient ( $C_{p_{rms}}$ ), histograms of surface pressures, power spectral densities (PSD) of surface pressures, and coherence analyses were performed to determine the required computational time to achieve the desired solution accuracy.

### **Results and Impact**

- Regions of elevated  $C_{p_{rms}}$  associated with high buffet loads were identified by the CFD, and match locations from wind tunnel tests.
- Comparisons of computational and experimental  $C_p$  histograms show similar trends, indicating the CFD is capturing the physical phenomena in the flow field.
- Comparisons of PSDs show that frequency content is similar, but not exact.
- CFD was used to investigate ability of alternate booster noses to reduce  $C_{p_{rms}}$  levels, with results corroborated by wind tunnel data.
- PSD analyses show that 0.5-0.8 second of simulated data is sufficient to resolve oscillatory forces on the SLS for the computed transonic cases.
- Coherence analyses show that 1 second of simulated data is necessary to correlate the unsteady pressures along the longitudinal axis of the SLS.

### **Role of High End Computing**

Computational predictions must be obtained in a timely manner, if the detailed information that CFD provides is to affect the vehicle design early in the program. The meshes required for DDES simulations of the SLS ranged from 35 million points to nearly 300 million points. The Fully Unstructured Navier-Stokes Three Dimensional (FUN3D) solver enables the acquisition of this data, through the use of 1000's of cores, without adversely affecting computational performance. Only systems such as the NASA Advanced Supercomputing (NAS) facility enable these types of computations, which have significantly contributed to the successful designs of NASA's SLS.

Presented at SC15 by:

Stephen Alter, NASA Langley Research Center, Stephen.J.Alter@NASA.Gov

Principal Investigator:

Greg Brauckmann, NASA Langley Research Center, Gregory.J.Brauckmann@NASA.Gov

Exhibit Theme: Space Exploration

## High-Fidelity Physics-Based Analysis and Design of Complex Configurations

### Overview

NASA missions require a wide range of modeling capabilities to design and analyze next generation flight vehicles with operating conditions spanning low subsonic speeds to high-velocity launch and reentry vehicles. In many cases, developing these new technologies requires improved understanding of fundamental physics beyond that embodied in analysis and design tools commonly deployed throughout the aerospace industry. The need to quantify the effects on important engineering design parameters caused by changing hundreds, or even thousands, of design parameters is also desirable. The ultimate goal is to use this technology in a formal optimization framework to systematically improve performance of the systems under development. To be successful, the simulation tools must use the most advanced numerical methods and physical models, sound mathematics, and good software design practices. The FUN3D toolset is actively under development at NASA Langley Research Center to meet these far-reaching agency objectives.

### Project Details

FUN3D is an unstructured-grid computational fluid dynamics suite of codes that includes high-fidelity fluid dynamics modeling, as well as a range of models for structural effects, acoustics, radiation, ablation, and chemistry. FUN3D has led in the development of simulation-based design tools and provides the world's foremost adjoint-based capability for accommodating large numbers of design variables. This FUN3D technology has also been instrumental in the development of adjoint-based methods for mesh adaptation and error estimation, which significantly reduces uncertainties in accuracy normally present in numerical simulations.

FUN3D has been scaled to 80,000 cores on meshes comprised of several billion elements. This capability is instrumental for successful analysis and design required for completion of many NASA missions.

### Results and Impact

FUN3D is widely used by numerous NASA programs, U.S. industry, the Department of Defense, and has been utilized by over 100 academic institutions for conducting research projects leading to advanced graduate degrees. Results from a previous collaboration with researchers at the Department of Energy were recognized with the Gordon Bell Prize, a prestigious award for outstanding achievements in high-performance computing. Because of the unique capabilities and performance on large-scale computing platforms, FUN3D continues to be widely used throughout the aerospace industry. Several recent and ongoing applications of FUN3D include:

- NASA programs including fixed-wing and rotary wing vehicles, sonic boom mitigation, and the development of NASA's new Space Launch System.
- Development of commercial crew spacecraft at companies such as SpaceX.
- Efficient green-energy concept development, such as distributed electric propulsion, wind turbine design, and drag reduction for long-haul trucking.

### Role of High-End Computing (Why HPC Matters)

FUN3D simulations require frequent solutions of several billion simultaneous equations. These intensive calculations can only be efficiently accomplished using high-end computing resources, such as the Pleiades supercomputer at the NASA Advanced Supercomputing facility.

Presented at SC15 by:

W. Kyle Anderson and Eric Nielsen, NASA Langley

Principal Investigator:

William.K.Anderson@nasa.gov, Eric.J.Nielsen@nasa.gov

Exhibit Theme: Aeronautics

## Computational Aero-Acoustics Open-Rotor Simulations for Green Aviation

### Overview

Reliable noise prediction capabilities are essential to enable novel fuel-efficient open rotor designs that can meet the community and cabin noise standards. Toward this end, both immersed boundary and structured curvilinear grid approaches have reached a level of maturity so that they are being frequently employed for specific real world applications within NASA. This work demonstrates that our higher-order methods provide the ability for aeroacoustic analysis of wake-dominated flow fields. This is the first of a kind aeroacoustic simulation employing an immersed boundary method for an open rotor application. In addition to discussing the peculiarities of applying the immersed boundary method to this moving boundary problem, we conducted a detailed aeroacoustic analysis of the noise generation mechanisms encountered in the open rotor flow. The simulation data (both immersed boundary and curvilinear based) was compared to available experimental data. The noise generation mechanisms were analyzed employing spectral analysis, proper orthogonal decomposition and the causality method.

### Project Details

There has been a renewed interest in pursuing contra-rotating open rotor propulsion technology due to the large potential of significantly reducing fuel assumption. A main concern for the design of such systems is that they must meet community noise and cabin noise standards. Hence, providing reliable noise prediction capabilities for contra-rotating open rotors is essential in the design of low-noise rotor propulsion systems. In recent years, NASA initiated several research efforts that have assessed the capabilities on noise source prediction and analysis of contra-rotating open rotors. Several low to high fidelity simulations tools were used to provide noise predictions and their results were compared to experimental data. One of the key challenges is to conduct efficient high-fidelity simulations of this relatively complex geometry.

### Results and Impact

In the present work, we independently employ higher-order accurate immersed boundary and structured curvilinear methods to simulate the highly complex flow field around an open rotor. Three key aspects were addressed in this work:

- (1) simulating the open rotor with our novel higher-order immersed boundary method,
- (2) simulating the open rotor with our boundary layer resolving body fitted structured curvilinear approach, and
- (3) detailed analysis of the most dominant noise generation mechanisms of the open rotor propulsion system.

In order to gain confidence in our simulation approach we conducted a detailed comparison with experimental data. Several post-processing tools, such as Fourier transforms, POD, and the causality method, were used to obtain a detailed understanding of the noise generation process. It was found that the large sensor plate contributed significantly to the broadband noise content due to interactions with the rotors and hub. Ongoing work includes assessing wind-tunnel wall effects.

### Role of High-End Computing

Using NASA's Pleiades supercomputer, the simulations consumed approximately three weeks of wall time on 500-2000 Intel Haswell cores. These simulations would not be possible without the use of a large cluster like Pleiades. The visualization team at NAS helped with generating ultra high quality renderings of the open-rotor flow physics.

Presented at SC15 by:

Michael Barad, NASA Ames Research Center, Michael.f.barad@nasa.gov

Principal Investigator:

Cetin Kiris, NASA Ames Research Center, cetin.c.kiris@nasa.gov

Exhibit Theme: Aeronautics

## **Improving High-Performance Computing Throughput with Memory Profiling**

### **Overview**

Analyzing per-thread memory usage of a high-performance computing (HPC) application is an important component in balancing memory load and improving performance. At the same time, controlling swap usage in HPC systems using memory swap is also necessary to help avoid network file system hangs. The Linux kernel 'cgroups' facility is a commonly used method of controlling per-user swap usage, but this approach does not provide any memory profile information. We describe an efficient technique for controlling swap space usage on a per-user basis, while providing profiling information to the user. The user has the option of generating plots with both memory usage per node and memory usage per thread.

### **Project Details**

We have developed a memory-monitoring tool that tracks a specific job associated with a user's process in the Slurm resource manager and monitors both the memory high-water mark per thread and the per node memory usage. The memory usage samples for each process ID are recorded in an XML file that can be used to produce graphs. This approach gives the user the opportunity to visually review the memory usage for a specific application. Using these graphs, the user can make the necessary code changes to improve memory distribution in computational processing threads.

The memory-monitoring tool runs in the background with low overhead, providing the opportunity for the user to control the swap ratio per machine used. In addition, we developed it in a low-level programming language using the native Slurm API functions; that way it is shell independent and can run on various architectures.

### **Results and Impact**

Our memory-monitoring tool provides the HPC user the following:

- The opportunity to control the memory swap of their run, thereby protecting the HPC network file system from hangs.
- The ability to monitor each thread's memory high-water mark.
- The ability to produce thread memory utilization graphs to help balance thread memory usage.
- The option to use the tool only when a job is initiated and not run it continuously.

### **Role of High-End Computing (Why HPC Matters)**

HPC provides scientists with the means for solving complex research and engineering problems. The responsible utilization of memory resources and the development of well-balanced HPC applications from the memory perspective improve the throughput of users' jobs and decrease expensive downtimes associated with out-of-memory network file system hangs.

Presented at SC15 by:

Georgios Britzolakis, NASA Goddard Space Flight Center, Georgios.Britzolakis-1@nasa.gov

Principal Investigator:

Laura Carriere, NASA Goddard Space Flight Center, laura.carriere@nasa.gov

Exhibit Theme: Supercomputing

## Unlocking the Secrets of Solar Storms

### Overview

Solar storms are large explosive events on the Sun, including flares and coronal mass ejections (CME), which are capable of ejecting billions of tons of million-degree magnetized plasma out into space. When impacting Earth, they can cause disruption in radio transmissions, damage to satellites, and massive damage to power transmission grids, leading to large-scale long-lasting power outages. Predicting such eruptions and their trajectories is a major goal of NASA's Heliophysics division under the Science Mission Directorate. Understanding the storms' underlying magnetic/plasma structure and evolution is vital for developing predictive capabilities. Our group has been engaged in several avenues of research towards this goal.

### Project Details

We use a sophisticated magnetohydrodynamic (MHD) model to improve our understanding of the initiation and propagation of solar storms. Our 3D MHD global coronal and heliospheric model has been used for over a decade to explore theoretical models of CME initiation, coronal heating, and to model the structure of active regions and the global corona and solar wind, including the inner heliosphere. The model allows us to perform CME "event studies," in which we use detailed observations as inputs, including observed magnetic fields. From the model we can produce simulated outputs that can be directly compared with observations, including white-light images as well as EUV and X-ray emission.

### Results and Impact

We have developed innovative methods for producing energized pre-eruptive magnetic field structures, together with the ability to trigger eruptions. We have also developed sophisticated post-processing techniques to analyze the topological structure of magnetic fields, which allow for a deeper understanding of the characteristics of solar eruptions. We have simulated in detail the "Bastille Day" event (July 14, 2000), which produced a fast CME exceeding 1700 km/sec, and a very large X5.7 flare. The simulations were able to reproduce some of the key observed features of this storm.

### Role of High-End Computing (Why HPC Matters)

Our MHD simulations of solar storms require grids of 50 million cells or more, and the integration of time-dependent equations for millions of cycles, necessitating the use of massively parallel computations. These are performed on the NAS Pleiades supercomputer. Typical simulations utilize thousands of processors and can take several days to complete. The critical importance of direct simulation results to the study of solar storms makes the use of HPC essential.

Presented at SC15 by:

Ronald Caplan, Predictive Science Inc., [caplanr@predsci.com](mailto:caplanr@predsci.com)

Principal Investigator:

Zoran Mikic, Predictive Science Inc., [mikicz@predsci.com](mailto:mikicz@predsci.com)

Exhibit Theme: The Universe

## Test-driving “The Sun in a Box”

### Overview

The Sun is the closest star to the Earth. It impacts us in a variety of ways, many of which are not well understood. For example, the Sun’s radiation plays a role in the Earth’s climate, whereas violent plasma storms on the Sun directly impact the space environment around Earth, which can lead to significant disruptions of our highly technology-dependent society (GPS reception, power grid stability). To understand how the Sun’s atmosphere is shaped and heated, and how it impacts Earth, we simulate a small piece of the Sun in a computational box that could hold the Earth six times, and compare our results with high-resolution observations of the Sun’s atmosphere from NASA’s recently launched Interface Region Imaging Spectrograph (IRIS).

### Project Details

We use advanced multi-dimensional radiative magnetohydrodynamic simulations to take into account the complex physical processes that power the Sun’s atmosphere. We perform detailed comparisons of our simulations with high-resolution observations of the low solar atmosphere from NASA’s IRIS small explorer satellite. This synergy of numerical modeling and observations is focused on revealing:

- How the outer solar atmosphere is heated to temperatures much in excess of the temperatures at the Sun’s surface through the dissipation of magnetic field energy
- How the complex interactions of magnetic fields, hydrodynamics and radiation fields shape the Sun’s atmosphere

### Results and Impact

We have compared images of the Sun from NASA missions with our simulations and have helped explain some of the puzzling findings of the IRIS mission. The simulations have helped reveal how small-scale magnetic fields impact the low solar atmosphere or chromosphere, and how friction between ionized and neutral particles dissipates magnetic energy and helps heat the chromosphere. The chromosphere is the source of most ultra-violet radiation that impacts Earth’s upper atmosphere.

Our simulations have not only provided critical insight into the physical mechanisms that drive some of the violent events observed with IRIS, but have also provided much needed advances in how to interpret the complex radiation that IRIS observes.

### Role of High-End Computing (Why HPC Matters)

Numerical modeling of the complex radiative transfer and physical processes, combined with the enormous contrasts of density, temperature and magnetic field within the solar atmosphere demands supercomputing resources. The simulations exploit the massively parallel capabilities of NASA’s Pleiades supercomputer. Approximately 2 million CPU hours (48 days on 1728 CPUs) are required to compute one hour of solar time.

### What’s Next

As we learn from discrepancies between observations and simulations, we are currently working on simulations with much higher spatial resolution to better resolve some of the dominant spatial scales in the solar atmosphere, as well as on larger computational volumes to capture the large-scale behavior of the Sun’s magnetic field. This is expected to lead to better agreement with IRIS observations.

Presented at SC15 by:

Mats Carlsson, University of Oslo, mats.carlsson@astro.uio.no

Principal Investigator:

Bart De Pontieu, Lockheed Martin Solar and Astrophysics Laboratory, bdp@lmsal.com

Exhibit Theme: The Universe



## **Navier-Stokes Simulation of Rotor Aeromechanics Using Adaptive Mesh Refinement**

### **Overview**

Helicopters, tiltrotor aircraft and drones provide many important civil and military services due to their ability to take off and land vertically without an airport. However, the spinning rotor blades produce tip vortices that may pass closely by other rotor blades when the vehicle is in forward flight. This blade vortex interaction (BVI) can create sharp pressure peaks on the blades with high sound levels, disturbing civilians below and alerting a foreign military of an approaching aircraft.

Dynamic stall is another important phenomenon that limits helicopter speed and performance while in forward flight. The lift generated by a spinning rotor blade as it advances into the headwind must be balanced as the blade retreats away from the headwind. Additional lift is generated on the retreating side by pitching the blade up to large angles of attack. However at some point the blade can stall with a sudden loss of lift. This occurs because a vortex forms at the blade leading, which then detaches and traverses toward the back of the blade. This also causes an undesirable nose-down pitching moment.

### **Project Details**

Both BVI and dynamic stall are difficult to accurately predict with Computational Fluid Dynamics (CFD), in part due to poor grid resolution on the rotor blades. NASA's Revolutionary Vertical Lift Technology (RVLT) project has successfully used adaptive mesh refinement (AMR) to locally refine the mesh in the rotor wake. Near-body AMR is used for the first time to locally refine the mesh on the rotor blades. Near-body AMR enables high-resolution accuracy of BVI pressure peaks and dynamic stall vortices that dynamically move throughout the rotor generated flow field.

### **Results and Impact**

High resolution near-body AMR is expected to greatly improve the accurate prediction of BVI sound levels and dynamic stall lift hysteresis for a spinning rotor. Since these phenomena dynamically move throughout the flow field, it is impractical to resolve the entire rotor blade to the levels required without AMR.

- Near-body AMR is applied to rotors in forward flight for the first time with the OVERFLOW Navier-Stokes CFD code.
- Near-body AMR is expected to verify if the remaining discrepancies between computation and experiment are due to aerodynamic resolution or deficiencies in the structural model of the flexible rotor blade.
- Improved spatial resolution is demonstrated for the challenging BVI and dynamic stall phenomena. This approach can lead to improved vehicle designs with higher flight speeds and reduced sound levels.

### **Role of High-End Computing**

These flow simulations were made possible by NASA's Pleiades supercomputer. Solutions typically required 1,024-2,048 cores for one week. Grid sizes vary from 80-800 million grid points, and solutions typically require 5-24 hours of wallclock time per rotor revolution, and up to 100TB of disk space in order to make movies of the unsteady flow simulations. Time-dependent animations are crucial in understanding what flow phenomena are responsible for the aerodynamic response.

### **What's Next**

BVI and dynamic stall are very important phenomena that are poorly resolved with traditional non-AMR methods. Further computations will be carried out to quantify and validate the approach with experimental measurements. Moreover, this method will be included in a Multidisciplinary Design Analysis and Optimization (MDAO) framework to reduce noise and control dynamic stall to achieve higher flight speeds.

Presented at SC15 by Principal Investigator:

Neal M. Chaderjian, NASA Ames Research Center, Neal.Chaderjian@nasa.gov

Exhibit Theme: Aeronautics

## How Will Exascale Data Change High-Performance Computing?

### Overview

Traditional high-performance computing (HPC) solutions are built to optimize large-scale simulations. At the NASA Center for Climate Simulation (NCCS) it is becoming increasingly common for these simulations to generate petabytes (PB) of data. Scientists are also attempting to combine observations and model outputs in new and novel ways. Since the datasets are so large and difficult to move, scientists typically employ the same HPC system that created the data to analyze that data, or they download subsets of data to their desktop environment. Neither of these solutions is optimized for such data access patterns and applications, leading to unbalanced usage of resources and extremely frustrated researchers.

### Project Details

Future exascale environments will bring together technologies that are optimized for both HPC and analytics on the same system. Exascale applications are expected to create exabytes (EB) of data. This situation will require performing analytics across all subsystems of the environment. Given the amount of power required to move data, future HPC environments must have the capability to perform *in-situ* analytics to minimize data movement. Until the HPC industry develops exascale environments, the NCCS is evolving its architecture toward this end state.

### Results and Impact

The NCCS is exploring hardware and software solutions from traditional HPC and emerging big data technologies, including:

- Advanced Data Analytics Platform (ADAPT) – Combines HPC and virtualization technologies to create a private cloud specifically designed for large-scale data analytics.
- Intel Head in the Clouds Project – Provides scientists with the ability to burst into commercial clouds as needed.
- Spatial Temporal Indexing – Provides a high-level spatial temporal indexing that uses the knowledge of the structure climate data file stored in Hadoop to dramatically speed up MapReduce jobs.
- Hadoop and Posix – Couples Hadoop and Posix file systems through the use of IBM's Global Parallel File System (GPFS) to enable traditional and emerging analytics on the same data.

### Role of High-End Computing (Why HPC Matters)

The Discover supercomputer (~80,000 cores with 33 PB of storage) is the primary compute capability within the NCCS and supports extremely high-resolution global climate model simulations and other projects. ADAPT is much smaller, containing only about 1,000 compute cores with approximately 5 PB of storage; it maintains a much higher storage-to-compute ratio than Discover. NASA scientists are using these systems to study our rapidly changing climate by tracking changes in surface water, vegetation, and biomass over different areas of the Earth.

### What's Next

The NCCS is building an infrastructure capable of employing these technologies in daily operations to bridge the gap between current petascale and future exascale environments. This infrastructure will combine storage and compute to enable Posix access to the data and the next generation of *in-situ*, storage-proximal analytics required for EB of data.

Presented at SC15 by Principal Investigator:

Daniel Duffy, NASA Goddard Space Flight Center, Daniel.Q.Duffy@nasa.gov

Exhibit Theme: Supercomputing

## Numerical Simulation of Slat Noise to Support Environmentally Responsible Aviation

### Overview

During takeoff and landing, the noise generated by high-lift devices on aircraft, such as the slat located at the front of the wing, cause large amplitude broadband sound waves. These waves propagate down to communities neighboring the airport causing disruption and annoyance. As part of the Environmentally Responsible Aviation (ERA) project at NASA, research on noise reduction concepts for airplanes is one of the key areas for next generation commercial aircraft. In order to reduce noise from high-lift devices, the underlying physics of the noise generation must first be understood. In an effort to better understand slat noise generation, both numerical simulation and experimental testing of a conventional high-lift slat device installed at high-angle of attack in an open jet facility have been performed.

### Project Details

The Launch Ascent and Vehicle Aerodynamics (LAVA) solver framework, developed at NASA Ames Research Center, was used for the high-fidelity time-accurate near-field CFD analysis. Noise generated by the slat is then propagated to the far-field using the LAVA Ffowcs Williams – Hawkings acoustic module. An experimental test was also performed in the Quiet Flow Facility (QFF) at NASA Langley Research Center to validate the computed results of the acoustic simulation. Time averaged surface  $C_p$  and both near-field and far-field PSD spectra are used for the validation.

### Results and Impact

Prior to the experimental test, LAVA CFD analysis was used exclusively to help design the experiment. Both steady and time-accurate RANS calculations were performed in both free-air and installed configurations for several angles of attack. During the experimental run, centerline  $C_p$  on the slat and main element were compared to the existing CFD results with excellent agreement.

To simulate the slat noise, a higher-order accurate finite-difference scheme for general non-orthogonal curvilinear coordinates was developed in LAVA along with advanced hybrid RANS/LES models. The increased accuracy in both the numerical discretization and the turbulence model allows fine scale structures to be modeled on relatively coarser meshes, reducing the computation resource requirements.

The impacts of the project on ERA include:

- Validation of the LAVA solver framework for predicting noise generation and sound propagation from high-lift devices
- A successful experimental test resulting from strong cross-center collaboration between NASA Ames and NASA Langley
- High-fidelity dataset to explore the noise generating mechanisms for a deeper understanding of slat noise.

### Role of High-End Computing (Why HPC Matters)

The high-fidelity time-accurate near-field CFD analysis was performed on the NASA Advanced Supercomputing (NAS) computer *Pleiades* using 500 cores and 62K core hours. For the acoustic propagation, the Fast-Fourier transform of long duration time series on both solid and permeable surfaces requires large amounts of memory. Large memory nodes on the *Pleiades* system containing upto 256 GB of memory are used to efficiently perform this part of the FWH procedure.

### What's Next

In addition to the conventional slat, a leading edge (LE) Krueger flap, which was designed using the LAVA CFD analysis tools, is currently being tested in the QFF tunnel. Acoustic simulations of the LE Krueger flap will be validated with the experimental data, and differences in the noise generation mechanisms between the slat and Krueger will be explored.

Presented at SC15 by:

Jeffrey Housman, NASA Ames Research Center, jeffrey.a.housman@nasa.gov

Principal Investigator:

Cetin Kiris, NASA Ames Research Center, cetin.c.kiris@nasa.gov

Exhibit Theme: Aeronautics

## **System-Level Assessment of Quiet Airframe Technologies**

### **Overview**

Diminishing the adverse environmental impact of aircraft noise on population centers near airports is a high priority for NASA's Aeronautics Research Mission Directorate. Airframe noise is a prominent component of aircraft noise during landing. Until now, system-level development and evaluation of viable airframe noise reduction technologies in a relevant environment could only be achieved with full-scale flight tests.

Researchers supporting the Environmentally Responsible Aviation (ERA) project are working to advance and accelerate the development of simulation-based, system-level, airframe noise prediction tools. Ultimately, the goal of creating such tools is to promote a paradigm shift in the design procedure from the current time-consuming and costly "cut-and-try" approach to a physics-based virtual design environment whereby the aeroacoustic evaluation of a noise reduction concept and its subsequent optimization can take place in an integrated fashion.

### **Project Details**

As continuation of our previous work, which was based on a full-span, full-scale baseline Gulfstream aircraft in landing configuration (with wing flaps deflected 39 degrees and the main landing gear deployed), we have simulated the same configuration with noise reduction (NR) technologies applied to the flaps and main landing gear. To determine the performance of the applied technologies and the ensuing noise abatement benefits, we designed our simulations to accurately resolve the local flow fields at three locations: the flap tips, the regions near the main landing gear, and the interaction zones between these two areas.

To date, we have completed coarse- and medium-resolution computations; fine-resolution results are in progress. We used Exa Corporation's PowerFLOW® Lattice-Boltzmann flow solver to perform these time-accurate simulations.

Depending on the geometric complexity of the NR technology being evaluated, our medium-resolution simulations required between two to four billion volumetric cells to discretize the space surrounding the aircraft surfaces.

### **Results and Impact**

Simulations of this kind present numerous challenges due to extreme geometric complexities and the wide range of spatio-temporal resolution that is required. Through a direct comparison with acoustic measurements obtained during a 2006 flight test, our previous results for the baseline full scale aircraft helped validate the accuracy of these computations. To the best of our knowledge, the simulations performed for this study constitute the first attempt to evaluate, on a system-level basis, airframe noise abatement technologies applied to a full-scale aircraft in landing configuration with most of the geometrical details included. Our study will help NASA foster a paradigm shift in the aircraft design process to include aeroacoustic requirements at the conceptual level, thus bringing closer the ultimate goal of confining aircraft noise within airport boundaries.

### **Role of High-End Computing (Why HPC Matters)**

Our full-scale aircraft simulations constitute the first step toward evaluating system-level aeroacoustic performance of a novel set of flap and landing gear noise reduction technologies on a full scale aircraft. As such, these complex simulations are very resource intensive. All computations for this study were performed on NASA's Pleiades supercomputer. Our medium-resolution simulations required approximately 4,000 to 8,000 processors and 275,000 to 2.0 million processor hours and generated 40-50 terabytes of data per case. Rendering and visualization of these large datasets were made possible by the NASA Advanced Supercomputing (NAS) Division's visualization team.

### **What's Next**

In the near term, we will apply the same gear noise reduction technologies to a full-scale Gulfstream G-III aircraft. Also evaluated on this aircraft will be additional concepts to reduce noise from landing gear cavities. These simulations will complement and guide the planned flight test of these concepts on NASA's G-III aircraft.

Presented at SC15 by Principal Investigator:

Mehdi R. Khorrami, NASA Langley Research Center, mehdi.r.khorrami@nasa.gov

Exhibit Theme:

Aeronautics

## **Star Formation From Giant Turbulent Molecular Clouds to Protostellar Clusters**

### **Overview**

The formation of stellar clusters and the high mass stars that form within them remains one of the most significant unsolved problems in all of astrophysics and is a major Grand Challenge. To attack such a challenge, simulations must include a large range of physical processes including gravity, supersonic turbulence, hydrodynamics, outflows, radiation (ionizing and non-ionizing), magnetic fields and chemistry. The high degree of non-linear coupling and feedback mechanisms among these processes and the extraordinarily large dynamic range in time and spatial scales of the problem add to the difficulty of such simulations. Many low mass stars are born in clusters containing massive stars, and Hubble Space Telescope observations show that disks around such stars are destroyed by the ionizing radiation from the stars themselves. Stellar clusters and massive stars thus lie at the center of the web of physical processes that shaped the Universe as we know it, yet the processes involved in their formation still remains elusive. Our work is closely connected to NASA's goals of determining the origin of stars and planets.

### **Project Details**

We use our state-of-the-art magneto-radiation-hydrodynamic 3-D adaptive mesh refinement (AMR) code ORION2 in these large-scale simulations, to follow the gravitational collapse over a dynamic range of over 32,000 of a highly turbulent, magnetized, massive molecular cloud down to the formation of multiple turbulent clumps forming inside these clouds whose further collapse results in the formation of star forming cores. The magnetized turbulent cores are then further evolved with AMR zoom-in simulations to the formation of stellar clusters. The huge dynamic range of the AMR simulations is effectively equivalent to  $3 \times 10^{13}$  volumetric cells. Our simulations include magnetic fields, radiation transport, turbulence, and highly energetic protostellar outflows.

### **Results and Impact**

Our simulations show that massive infrared dark cloud filaments (IRDCs) form by gravitational collapse in a strongly magnetized, highly supersonic turbulent environment with a column density that has a lognormal-like probability distribution function at early stages of collapse evolving into a power law later. Proto-stars inside dense cores have a proto-stellar mass function well explained by the theoretical two-component turbulent core model at all times. After 700,000 years, the distribution matches very well with the Chabrier stellar initial mass function and the multiplicity of binary, triple, and quadruple proto-stellar systems formed in the simulation match well with observations.

### **Role of High-End Computing (Why HPC Matters)**

Our simulations were performed on NASA's Pleiades supercomputer with each simulation using 1000-4000 processors requiring over 6 million cpu hours to complete over several months of computation. Simulations of this enormous range of scale and highly coupled physics, producing about a hundred Terabytes of data, and follow up data visualization would not be possible without NASA's vast computational resources.

Presented at SC15 by Principal Investigator:

Richard Klein, University of California, Berkeley and Lawrence Livermore National Laboratory,  
rklein@astron.berkeley.edu

Exhibit Theme:

The Universe

## **The Application Performance and Productivity Group**

### **Overview**

It is amazing how many things must go “just right” in order to do calculations optimally on modern high performance computing systems. The layout of memory and the ordering of arithmetic operations can greatly affect the movement of data, which typically has a significant effect on performance. Throwing more compute resources at the problem can often speed things up, but leads to a question of what is meant by “optimization”. Are we minimizing time to solution or use of resources? Once that is answered, other questions emerge. How do we go about the process of optimization? What do different performance analysis tools bring to the process?

### **Project Details**

NASA's High End Computing Capability Project has an Application Performance and Productivity (APP) team that provides a range of consulting services for users. In addition to helping port codes or find causes for job failures, the group also works closely with users to understand and improve code performance. When a user asks the APP team for help improving code performance, it starts by acquiring the code, a representative dataset, and expected results. Then begins an iterative process to identify performance issues (hotspots) that we attempt to reduce or eliminate. Performance analysis tools are essential to help in this process. For example, identification of bottlenecks is facilitated when tools report things such as:

- where the code is spending walltime,
- if a hotspot has been compiled using scalar or vector instructions, and
- when there is an imbalance between different threads of execution.

Each performance analysis tool has its own strengths and weaknesses. As a consequence, the APP team uses a variety of performance analysis tools during the optimization process. In the booth presentation we will discuss five different tools and how each can contribute.

### **Results and Impact**

In the past year, HECC users have asked performance consultants on the Application Performance and Productivity team for help with optimizing their codes and the team has responded with a number of notable performance successes:

- SFILTER – more than a factor of 2x speedup
- VULCAN – 45% faster code and knowledge transfer
- MULTIA – 4x faster code and more to come
- Fortran CUDA wrapper – optimization through CUDA libraries

In the case of sfilter, the improvements allowed the science team to perform a simulation that previously would not run. Furthermore, the modified code with new Intel compiler and job partitioning resulted in a savings of more than 2.47 million SBUs. 1 SBU is equivalent to one hour on a Westmere node.

The VULCAN effort resulted in a code that ran 46% faster than before. Aside from the ability to complete more simulations, this effort was part of an ongoing collaboration between NASA/Ames and NASA/Langley to share knowledge regarding ways to improve application performance.

The MULTIA effort lasted about three weeks and resulted in a 4x speedup. The code owner really needs another 25x in order to achieve his science goals.

The CUDA wrapper makes a subset of the existing, highly optimized CUDA libraries accessible to Fortran programmers

### **What's Next**

In addition to investigating codes at user request, the APP team will continue collaborating with other NASA centers to share knowledge. The team also works closely with the Pleiades system administrators to identify user codes that trigger instabilities within the system.

Presented at SC15 by:

Robert Hood, NASA Ames Research Center, [Robert.hood@nasa.gov](mailto:Robert.hood@nasa.gov)

Principal Investigator: Dan Kokron, NASA Goddard Space Flight Center, [daniel.s.kokron@nasa.gov](mailto:daniel.s.kokron@nasa.gov)

Exhibit Theme: Supercomputing

## Climate Data Downscaling Projects at NASA

### Overview

It is common that we use global climate models or Earth system models to perform climate projection into the future. Because of the long integration time and the tremendous computing resources required for such a projection, the model resolution is typically not at a spatial scale fine enough for climate assessment or decision support purposes. The policy and management communities are hungry for climate projection information at finer resolution for decision-making purpose. For this reason, a number of “downscaling technologies” have been developed over the years to bring the climate projection information to the local level.

With the consequence of climate change beginning to emerge, there is a continuous need for better quantifying the quality of downscaled climate projections. In this talk I will give an overview on NASA’s efforts to produce high-resolution data sets and to understand the various techniques, the limitations including the risks of using these techniques, and finally, I will provide a view on possible future researches in this area.

### Project Details

In the past couple of years, NASA supported a number of regional to local climate projection activities: NASA Climate Adaption Science Investigators focused on climate resilience at NASA center level, National Climate Assessment (NCA) Capacity Building focused on data sets and tools to support NCA, NCA Indicators focused on creating simple indicators specifically designed for decision support, Assessing the Fidelity of Dynamical Downscaling with the NASA Unifies-WRF Model focused on understanding the credibility of dynamical downscaling technique using a regional climate model. All of these projects have a component in creating or using downscaled climate information. We will show some of the results produced in 2015. We will also discuss a possible way forward.

### Results and Impact

Multiple downscaled products had been generated using different techniques (e.g. statistical, dynamical downscaling, data assimilation). These data products are supporting the recent National Climate Assessment and will be used in the future assessments.

### Role of High-End Computing (Why HPC Matters)

Long-term (typically at least 100-year) climate projections require long (wall clock) model integration on a supercomputer. Because of significant computing resource needed for the simulation and the associated post processing, the models are typically configured to use lower than optimal spatial resolution. HPC systems have been used to support the model integration. HPC are also used for model output analysis, model and observational data assimilation, and high resolution data product generation.

Presented at SC15 by Principal Investigator:  
Tsengdar Lee, NASA Headquarters, [tsengdar.j.lee@nasa.gov](mailto:tsengdar.j.lee@nasa.gov)

Exhibit Theme: Our Planet

## Taking the Dual Path to AMR Visualizations

### Overview:

Adaptive Mesh Refinement (AMR) techniques are rapidly growing in use for a wide range of Computational Fluid Dynamics (CFD) problems. With the increasing use of AMR solver codes comes the need for visualization and analysis tools that can efficiently process the results from those simulations. We have developed an approach based on working with the dual of the AMR mesh hierarchy in order to achieve continuous interpolation and visualization results comparable to those based on traditional meshing schemes. Unlike some similar techniques, our approach works even in cases where the AMR data have refinement ratios other than 2. We have applied our technique to data from a variety of CFD solvers, including CART3D, PowerFLOW, LAVA and ENZO. Those solvers are producing results key to achieving many NASA mission objectives

### Project Details:

The rapid growth in the use of AMR-based flow solvers has created new challenges in CFD visualization. Initially, we used reformatting and resampling techniques in order to convert the AMR data into a form that would work with existing tools. This was not a satisfactory long-term solution because the conversions introduced additional demands on HEC resources and slowed the analysis process. In particular, when analyzing results from simulations that were already pushing the envelope in terms of what the system was capable of, we wanted analyses techniques that were also capable of handling such demanding data without exceeding the limits of the system. In response, we developed a dual-mesh technique for AMR data that can work with the data “as is” and yet still produce visualizations comparable to those based on more traditional meshing schemes. Our approach does not require significant pre-processing or additional memory usage. Initially, we developed our technique to work with galaxyformation data from ENZO, a hydrodynamic astrophysics simulation code. Over time we have ported our technique to work with data from two in-house codes — CART3D and LAVA — and one commercially-available code: PowerFLOW.

### Results and Impact:

Our visualizations based on our dual-mesh analysis techniques have been featured by computational scientists in a variety of venues, including conference presentations, academic publications and industry journals such as Aviation Week. We anticipate that AMR-based codes will continue to account for a significant portion of the usage of the Pleiades HEC system, and thus the demand for visualizations of AMR data will remain robust.

In the case of the analysis of CART3D data, we have developed stand-alone applications that enable project scientists to quickly visualize their data and produce presentation-quality results on their own. One application in particular, raycart, produces ray-cast volume-rendering images, including images similar to those produced by experimental Schlieren techniques. This application has been enthusiastically adopted for use in support of a variety of projects, including analyses related to the Space Launch System (SLS) development and the modeling of meteors.

### Role of High-End Computing (Why HPC Matters)

Many of the studies conducted on supercomputing systems such as Pleiades are only feasible on large HPC systems. These studies tend to push the boundaries of what the system is capable of, not only in processing, but also memory and network capability. We want visualization techniques and tools that are also suited to such a demanding environment. We believe our dual-mesh visualization approach effectively meets those challenges.

Presented at SC15 by:

Patrick Moran, NASA Ames Research Center, [patrick.moran@nasa.gov](mailto:patrick.moran@nasa.gov)

Exhibit Theme: Supercomputing



## Liquid Rocket Engine Feed Line/Engine Interface Flow Uniformity Assessment

### Overview

The propulsion system of the Space Launch System (SLS) requires the integration of feed lines with turbomachinery. From a structural perspective there are many differences between these elements, but from the perspective of the propellant there is no distinct boundary between the two, and flow environments generated in the feed lines or the turbopump can interact and affect each other. Mission requirements dictate the geometric and operational configuration of the feed lines and pumps. This work is an assessment of the flow environment in the SLS liquid hydrogen (LH) feed line, upstream of the RS-25 low pressure fuel pump (LPFP) at nominal SLS operating conditions

### Project Details

Three simulation cases were run using Loci/CHEM with full scale geometry and liquid hydrogen as the propellant. 1) SLS LH feed line with no angulation and no inducer 2) SLS LH feed line with no angulation, coupled to RS-25 LPFP, 3) SLS LH feed line with 4.5 degrees gimbal angle, coupled to RS-25 LPFP.

The computational domain includes LH duct bend, ball-strut tie rod joint, a gimbal joint, and the RS-25 LPFP.

Operating conditions (i.e. LH temperature, pressure flow rate, pump speed) are for SLS nominal operation RS-25 109% rated power level. Simulations were run with single phase calculations (no cavitation).

Flow uniformity was calculated for the three cases and compared to each other. The various sources of flow variation were identified and their relative contribution to the flow field was determined.

### Results and Impact

While the LH duct geometry and components do generate a strictly non-uniform flow field, the LPFP is the largest contributor to flow non-uniformity near the feed line/engine interface. This is due to the proximity of the LPFP to the interface location and the swirling back flow through the blade tip gaps near the duct walls. Duct angulation at the gimbal joint may have a small effect on the pump back flow environment.

### Role of High-End Computing (Why HPC Matters)

Modeling the full scale geometry with fine details (e.g. tight LPFP tip clearances) requires a large number of elements (400M+), which is only feasible to run with the High-End Computing resources.

### What's Next

Future work will entail an examination of forces and moments on LPFP blades as well as assessing any changes in pump steady performance due to duct angulation.

Presented at SC15 by Principal Investigator:

Andrew Mulder, NASA Marshall Space Flight Center, [andrew.d.mulder@nasa.gov](mailto:andrew.d.mulder@nasa.gov)

Exhibit Theme: Space Exploration

## **Sensitivity Analysis for Chaotic Fluid Simulations**

### **Overview**

Formal sensitivity analysis methods play a critical role in computational science, enabling fields such as design optimization, uncertainty quantification, and error estimation. Such techniques provide a rigorous measure of how simulation outputs vary with changes in simulation inputs. For example, in the field of computational fluid dynamics, an aircraft designer can benefit greatly from knowing how vehicle performance varies with wing shape or how noise in the vicinity of an airport is affected by a specific landing gear configuration.

Traditional sensitivity analysis approaches are routinely used for sensitivities of time-averaged quantities in flows governed by the unsteady Reynolds-averaged Navier-Stokes (RANS) equations. However, as today's leading-edge simulations begin to adopt higher-fidelity eddy-resolving techniques such as Large Eddy Simulations (LES), Hybrid RANS-LES approaches, and Direct Numerical Simulations (DNS), classical sensitivity analysis methodologies break down, whether applied in forward or reverse (adjoint) mode. In these applications where turbulent fluctuations are being directly simulated, the flow fields and associated outputs such as lift and drag are chaotic in nature. Traditional sensitivity analysis procedures fail to provide useful information here because of the well-known "Butterfly effect", namely the high sensitivity of simulation outputs to initial conditions for chaotic dynamical systems.

### **Project Details**

In this effort, the NASA Langley FUN3D development team partners with researchers at Massachusetts Institute of Technology (MIT) and the National Institute of Aerospace to implement and evaluate a new approach developed at MIT for sensitivity analysis of chaotic simulations known as Least Squares Shadowing (LSS). LSS computes sensitivities using the "shadow trajectory", a phase space trajectory for which perturbations to the flow field do not grow exponentially in time. To efficiently compute many sensitivities for one objective function, we use an adjoint version of LSS.

### **Results and Impact**

The approach has been evaluated for flow over a two-dimensional NACA 0012 airfoil at high angle of attack, for which chaotic shedding from the upper surface of the wing occurs. Using the time-averaged lift as a representative output function, the solution of the traditional unsteady discrete adjoint equations grows exponentially in time. To overcome this barrier, the team has been able to successfully apply the LSS methodology to compute a bounded discrete adjoint solution for such a flow for the first time.

If the LSS approach can be successfully extended to large-scale three-dimensional simulations, a long-standing barrier for the sensitivity analysis community will have been removed. The methodology would have a broad impact on fluid simulations of relevance to NASA across the speed range. And since chaotic dynamics can occur in other disciplines such as aeroelasticity, the potential benefits of the approach reach far beyond just the fluid dynamics community.

### **Role of High-End Computing (Why HPC Matters)**

The challenge in applying LSS lies in its extreme computational requirements. Rather than the traditional initial value problem solved in everyday computational fluid dynamics, the LSS formulation requires the solution of a globally coupled space-time system of equations. For the simple airfoil test case described above, a coarse spatial grid consisting of just 100,000 grid points is used, with 15,000 time steps. While the traditional CFD simulation for this problem could be solved in minutes on a reasonable desktop machine, the LSS formulation results in a system of billions of coupled equations requiring a parallel decomposition in time across 15,000 processors and many terabytes of input data generated by FUN3D.

Simulations currently being performed at NASA are at least six orders of magnitude larger than the baseline problem studied here. The LSS methodology is an ideal candidate for leveraging the next generation of exascale-class hardware.

Presented at SC15 by:

Eric Nielsen, NASA Langley Research Center, [Eric.J.Nielsen@nasa.gov](mailto:Eric.J.Nielsen@nasa.gov)

Principal Investigator:

Qiqi Wang, Massachusetts Institute of Technology, [qiqi@mit.edu](mailto:qiqi@mit.edu)

Exhibit Theme: Aeronautics

## **Transforming the Passenger Aircraft for Fuel Efficiency**

### **Overview**

The next advance in aircraft technology may come from how the propulsion system is integrated into the aircraft. Traditionally, the engines hang under the wing or on the side of the rear fuselage. In all cases, great care is exercised to avoid sucking the boundary layer (slower air next to the aircraft) into the engine to assure minimum flow distortion. If a distortion-tolerant fan can be developed, game-changing methods of integrating the engine into the fuselage are possible. The present project investigates if there is an aerodynamic benefit associated with integrating the engine so that it is ingesting the boundary layer.

### **Project Details**

To quantify the benefit of boundary layer ingestion (BLI), an aircraft design developed at MIT is investigated side-by-side with a conventional engine configuration. Using computational fluid dynamics (CFD), both the BLI configuration and the conventional configuration are modeled. The surface of the aircraft and the air around it is modeled using 166 million points with approximately 70 grids that overlap each other's volumes. The flow of air is then solved on the Pleiades supercomputer at NASA Ames using mathematical equations that describe and predict fluid flow. The resulting airflow is analyzed to assess the power required to fly at cruise condition. A comparison of the power requirements results in quantification of the benefit achieved by using BLI engine integration.

### **Results and Impact**

The computational results show that at low speeds (70 to 100 mph), the BLI configuration has a substantial benefit over a conventional configuration. If this benefit can be shown to exist at higher speeds and distortion-tolerant fans become reality, all aircraft may be able to save a substantial amount of fuel for each flight; resulting in substantial savings for all operators of aircraft. It would also result in lower flight costs for airlines, possibly translating to lower ticket prices for passengers. The computational results also show that the loss associated with the fuselage boundary layer is translated into more efficient fan operation for the BLI configuration and is the cause of the benefit.

### **Role of High-End Computing (Why HPC Matters)**

Accurate CFD simulations require a large number of grid points to capture the physics of the airflow along the aircraft surface and in its vicinity. A grid for one half of the BLI configuration, for example, consists of approximately 166 million grid cells, which require a large amount of computational power to solve. A single analysis run of the aircraft and the domain around it at a single flight condition and one engine power setting requires approximately 25,000 processor-hours on the latest incarnation of the Pleiades supercomputer at NASA Ames Research Center. The geographically diverse team working on both coasts of the United States also takes advantage of the fast networks for efficient collaboration.

### **What's Next**

The benefit of the BLI engine integration must now be shown to exist at higher speeds at which commercial aircraft fly. The next step is to redesign the aircraft so that it is efficient at the higher speeds and then assess the BLI benefit at  $M=0.7$  to  $M=0.85$  to do trade studies of fuel efficiency vs. speed.

Presented at SC15 by Principal Investigator:

Shishir Pandya, NASA Ames Research Center, [Shishir.pandya@nasa.gov](mailto:Shishir.pandya@nasa.gov)

Exhibit Theme: Aeronautics

## NASA's Exploration in Quantum Computing

### Overview

NASA's quantum computing project is a long-term experiment to assess the potential of quantum computers to perform computations that are difficult or impossible using conventional supercomputers in a realistic timeframe. The project is a collaboration among teams at NASA, Google, and the Universities Space Research Association (USRA).

### Project Details

NASA researchers are employing a 1097-qubit D-Wave 2X system to investigate areas in which quantum computing and quantum algorithms may dramatically improve our ability to solve difficult computational challenges in aeronautics, Earth and space sciences, and space exploration. Current studies focus on quantum approaches to NASA optimization problems in applications such as air traffic control; machine autonomy, fault diagnosis, and robust system design; and mission planning and scheduling.

The team's work has two main thrusts. Its exploration of quantum annealing applications provides insight into which application area have the greatest potential for impact from quantum annealing technologies, advances programming techniques for quantum annealers, and provides guidance for the design of future quantum annealers. As its second thrust, it analyzes quantum annealing theoretically and numerically to better understand when it provides enhancement and through what mechanisms.

### Results and Impact

Key outcomes from the preliminary research studies completed in 2013-2015 include:

- Established benchmark sets of problems suitable for near-term quantum annealers from a variety of optimization problems in, planning, scheduling, fault-diagnostics, graph analysis, transport networks, and machine learning.
- Established best practice techniques for programming a quantum annealer, including determining the best way to "compile" fully-connected subproblems into an annealer.
- Developed novel hardware characterization and calibration techniques that increase performance 10 times or more by mitigating intrinsic analog control errors present in annealing devices.
- First examples of quantum annealing using uniquely quantum effects to outperform classical annealing in realistic systems subject to noisy conditions.

### Role of High-End Computing (Why HPC Matters)

NASA is exploring the potential of operating quantum systems as special-purpose processors attached to classical supercomputers located at the NASA Advanced Supercomputing facility at Ames Research Center. We are excited about possibility of tackling more efficiently the hardest computational problems through this integration.

### What's Next

Our researchers are collaborating with teams around the world who are building quantum hardware, providing design insights and hardware characterization. Scientists expect that quantum computing has the potential to vastly improve performance on a wide range of tasks, leading to new discoveries and technologies and significantly changing the way we solve challenging real-world problems.

Presented at SC15 by:

Alejandro Perdomo-Ortiz, NASA Ames Research Center, [alejandro.perdomoortiz@nasa.gov](mailto:alejandro.perdomoortiz@nasa.gov)

Principal Investigator:

Andre Petukhov, NASA Ames Research Center, [andre.g.petukhov@nasa.gov](mailto:andre.g.petukhov@nasa.gov)

Exhibit Theme: Supercomputing

## **Projecting Sea Level Rise by Modeling the Evolution of Ice Sheets**

### **Overview**

Being able to project sea-level rise (SLR) in a changing climate is of paramount importance if we want to be able to mitigate its societal and economical impact. One of the most significant contributions to SLR comes from fresh-water fluxes originating in the melting polar ice sheets of Greenland and Antarctica. Understanding and projecting the evolution of both these ice caps is therefore a priority, which NASA is involved in through ongoing support for the Ice Sheet System Model (ISSM), at Jet Propulsion Laboratory (JPL) and the University of California at Irvine (UCI) developed modeling software capable of assimilating remote sensing data into projections of the mass balance of Greenland and Antarctica. The goal is to accurately assess the future contribution of fresh-water fluxes to the ocean, and to understand how this contribution will impact SLR in the coming decades.

### **Project Details**

Accurately modeling the evolution of ice sheets not only involves complex physical processes that are still not fully understood, but also a significant amount of computational resources. Furthermore, a large amount of input data needs to be processed that come from various sources, such as regional climate models. To address these issues ISSM, which is a state-of-the-art finite element model, takes full advantage of a wide range of high performance computing environments. This is essential in order to deal with the high number of degrees of freedom that arise as a result of the finite element discretization process and the Stokes's equations that govern ice flow.

### **Results and Impact**

The results of the ISSM team's efforts have had a widespread impact on the Cryosphere community. Among its most significant contributions are the ability to automatically differentiate the ISSM code base in order to assimilate NASA's remote sensing data, as well as providing a means of performing uncertainty quantification. Moreover, optimization procedures utilizing NASA's Operation IceBridge data have lead to the an accurate and physically consistent bedrock topography of the Greenland ice sheet. Finally, the ISSM team and its users are constantly pushing the envelope by improving the modeling of physical processes such as calving front dynamics, ice/ocean interactions, hydrology at the ice/bed interface and paleo reconstructions of polar ice caps spanning thousands of years, with the goal of improving model spin-ups.

### **Role of High-End Computing (Why HPC Matters)**

Access to NASA supercomputing resources is an indispensable asset to the ISSM team. By leveraging these resources scientists using ISSM can complete jobs what would take weeks on conventional computing resources into a matter of hours. For large runs involving thousands of nodes with large amounts of inter-process communication, there is no other viable alternative. Of particular importance to us is the Pleiades Supercomputer at NASA's Ames Research Center, which we make substantial use of on a regular basis.

Presented at SC15 by:

Gilberto Perez, University of California, Irvine, gilberp@uci.edu

Principal Investigator:

Eric Larour, Jet Propulsion Laboratory, eric.larour@jpl.nasa.gov

Exhibit Theme: Our Planet

## Feasibility of Quantum Annealing in Earth Sciences: Data Assimilation and Image Registration

### Overview

Yuri Manin, Richard Feynman, and David Deutsch first conceptualized the idea of a quantum computer in the early 1980's. Quantum computers leverage quantum mechanical effects, such as superposition and entanglement, to efficiently solve computational problems. Jeffrey Fahri and co-authors first proposed the idea of Quantum Adiabatic Computing (QAC) in 2000 --- quantum computing through adiabatic evolution. QAC can be used to solve optimization problems, a task required in many scientific fields including Earth Sciences.

D-Wave Systems, Inc. made available the first commercial QAC system in 2013. In the past year, NASA's Quantum Artificial Intelligence Lab has acquired a 1000+ qubit (quantum bit) QAC system where research is actively being carried out. The D-Wave QAC computer is a physical Ising system and is capable of solving quadratic binary unconstrained optimization (QUBO) problems of the form

$$H(q) = \sum_{i,j} a_{i,j} q_i q_j.$$

To learn and assess the feasibility of QAC in areas of Earth Sciences, we are investigating its applicability in performing data assimilation and image registration to generate satellite data products.

### Project Details

Remote sensing instruments like OCO-2 (and AIRS and CRIS) are able to infer concentrations of CO<sub>2</sub> in atmospheric profiles, however they lack the sensitivity near the surface to provide *direct* estimates of sources and sinks. Nevertheless, these observations do contain information about terrestrial CO<sub>2</sub> fluxes, and one way to access this information is through inversion algorithms like data assimilation.

The collection of multi-temporal data to perform the assimilation is only possible through accurate calibration using image registration. Both image registration and data assimilation require minimization of real valued objective functions. In this project, we investigate different approaches of casting these objective functions into a QUBO and solving them using the D-Wave QAC computer.

For data assimilation, we are investigating two methods: (1) A polynomial approximation cast into a QUBO and (2) regression using a Hopfield Neural Network (HNN). In (1), real valued parameters are cast into binary form using fixed precision and higher order terms are reduced by introducing ancilla variables. In (2), the training of the HNN is performed on the D-Wave system. To train the HNN requires the minimization of a QUBO for each datum in the training set. In the case of image registration, we are investigating the two methods: (1) the construction of a conflict graph between selected features of the two images being registered and (2) regression using a HNN.

### Results and Impact

Currently we are in the beginning of this two-year project and the latest results will be presented.

### Role of High-End Computing (Why HPC Matters)

Earth Science applications require large computational resources. Areas such as weather prediction have a strict time-to-solution requirement, and the available HPC resources in part determine the prediction accuracy. Quantum computing offers a drastically different approach to HPC and has the potential to play an important role in Earth Sciences.

Presented at SC15 by:

Craig Pelissier, NASA Goddard Space Flight Center, [craig.s.pelissier@nasa.gov](mailto:craig.s.pelissier@nasa.gov)

Principal Investigator:

Milton Halem, University of Maryland, Baltimore County, [halem@umbc.edu](mailto:halem@umbc.edu)

Exhibit Theme: Supercomputing

## **Tripling the Compute Capacity at the NASA Center for Climate Simulation**

### **Overview**

The NASA Center for Climate Simulation (NCCS) provides high-performance computing (HPC) and data storage services for NASA's Science Mission Directorate. As researchers increase the resolution of their weather and climate models, supercomputers and data storage systems must rapidly grow to support expanding computational requirements and the petabytes of data being generated.

### **Project Details**

With a Scalable Compute Unit (SCU) design, the NCCS Discover supercomputer has been growing at a rate of approximately one SCU addition per year since 2006. In the fall of 2014, the NCCS embarked on an unprecedented expansion plan for Discover that would triple the cluster's computational capacity and double its online disk capacity. This effort also required the decommissioning and de-installation of five older SCUs and six aging storage systems.

### **Results and Impact**

Beginning in late November 2014, the NCCS began installing three SCUs that have increased the overall size of the Discover cluster from 43,000 processor cores performing 1.1 petaflops peak to 80,000 processor cores performing 3.5 petaflops peak. Staff installed these three SCUs in rapid succession over a period of only 6 months—the fastest and largest expansion ever executed at the NCCS:

- SCU10 – 30,240 Xeon Haswell cores, 1.2 petaflops peak (November 12, 2014)
- SCU11 – 17,136 Xeon Haswell cores, 713 teraflops peak (December 15, 2014)
- SCU12 – 17,136 Xeon Haswell cores, 713 teraflops peak (May 26, 2015)
- Total System – 80,000 Xeon (Sandy Bridge and Haswell) cores, 3.5 petaflops peak

To handle the large amounts of data being produced by this compute expansion, the NCCS has also integrated an additional 20 petabytes of disk storage into the Discover environment, bringing the total capacity to 33 petabytes.

### **Role of High-End Computing (Why HPC Matters)**

As a result of this large system upgrade, the Global Modeling and Assimilation Office (GMAO) at NASA Goddard Space Flight Center has been able to use 30,000 Haswell cores to run a global weather simulation at 1.5-kilometer resolution with more than 200 million grid cells. This represents the highest-resolution global simulation to date and the first ever to include interactive aerosols and carbon. The GMAO is exploiting these expanded NCCS resources to increase the resolution of their operational models and to examine features in the models that could not be resolved at lower resolutions.

Presented at SC15 by Principal Investigator:

Bruce Pfaff, NASA Goddard Space Flight Center, [bruce.e.pfaff@nasa.gov](mailto:bruce.e.pfaff@nasa.gov)

Exhibit Theme: Supercomputing

## Advanced Data Analytics Platform at the NASA Center for Climate Simulation

### Overview

The Advanced Data Analytics Platform (ADAPT) is a Platform-as-a-Service (PaaS), managed Virtual Machine (VM) environment providing compute, storage, and data management as well as user data publication services. Its charter is to streamline the process of putting flexible computational tools at the disposal of NASA scientists, thus allowing them to focus on Earth science rather than computer science. Currently a system administration- intensive framework, ADAPT is evolving into a more cloud-like venture operating under OpenStack.

### Project Details

ADAPT's feature set includes:

- A mix of physical and virtualized high-performance computing (HPC), storage, and network assets linked by QDR/FDR InfiniBand (IB) networks.
- VMs customized for data-intensive science applications and flexible for various loads, applications, and operating systems (Linux and Windows).
- iRODS, Ramadda, and ArcGIS for distributed data management.

### Results and Impact

ADAPT is currently being leveraged to better understand the impacts of a rapidly changing climate in Alaska and western Canada. In partnership with NASA's Terrestrial Ecology Program, the NASA Center for Climate Simulation is supporting the Arctic Boreal Vulnerability Experiment (ABOVE) field campaign. ABOVE will challenge scientists to take measurements in the field, study remote observations, and then run computational models to verify and project trends.

ADAPT has also been used in engineering applications such as simulating near-Earth object (NEO) discovery to assess options for improving system fidelity and rapid chip design simulations for the Laser Communications Relay Demonstration.

### Role of High-End Computing Resources

At the center of the ADAPT architecture is a multi-petabyte, open-source Gluster storage environment. Surrounding the storage with multi-core, large-memory, IB-interconnected HPC nodes facilitates:

- Ephemeral VMs for 'as needed' processing and persistent VMs for ongoing data delivery services.
- Highly available databases.
- Server-side graphical processing and rendering of data.
- Leveraging of large, shared datasets.

Scientists interact with the ADAPT team to provision resources and launch ephemeral VMs. The data-centric, virtual system approach significantly lowers the barriers and risks to organizations that require access to HPC solutions on an on-demand basis. As OpenStack is introduced into the environment, scientists will become more self-sufficient in deploying VMs.

### What's Next

Continued ADAPT evolution will concentrate on broadening the application and user footprint while augmenting both computational storage resources. OpenStack will bring a new level of functionality and flexibility to the platform.

Presented at SC15 by:

Daniel Pierce, Hoot Thompson, Jonathan Mills, NASA Goddard Space Flight Center  
daniel.l.pierce@nasa.gov, hoot.thompson@nasa.gov, jonathan.mills@nasa.gov

Exhibit Theme: Supercomputing



## Discovering How Galaxies Form by Comparing Simulations with Hubble Images

### Overview

Hubble Space Telescope images reveal that when galaxies are first forming they are very different from the mature galaxies that we see nearby. New galaxies are sausage-shaped, and a majority of star-forming galaxies have giant clumps of stars hundreds of times more massive than in nearby galaxies. We used to think that galaxies grow in size as they grow in mass, but observations show instead that galaxies typically undergo one or several compaction processes that rapidly shrink their sizes. We have seen all of these phenomena in our latest high-resolution supercomputer simulations, and this helps us understand how and why these processes are occurring.

### Project Details

We live in the disk of our Milky Way galaxy, which has a spheroid at its center. Other nearby large galaxies are also combinations of stellar disks and spheroids. But observations, especially from Hubble, show that galaxies in the early universe are very different. Understanding this is a challenge for theoretical astrophysicists.

Until a few years ago, even the best galaxy simulations were not much like real galaxies. But now increasing supercomputer power and improved simulation codes are producing galaxies that closely resemble the images and spectra of galaxies both nearby and in the early universe.

### Results and Impact

The key improvements in galaxy simulations have been starting from realistic cosmological initial conditions, increasing the resolution, and improving the realism of the simulated astrophysics. In nearby regions where stars are forming we see “pillars of creation” – regions where dust protects forming planetary systems from the intense radiation of the most massive short-lived stars. Our simulations now include such radiation, as well as the effects when these massive stars explode as supernovas.

Our galaxy simulations support the largest-ever Hubble Space Telescope project, the Cosmic Assembly Near-infrared Deep Extragalactic Legacy Survey (CANDELS). They explain many of the key features observed in galaxies that are in the process of formation, and make predictions for subsequent telescopes such as James Web Space Telescope that will provide more detailed images and spectra of even more distant galaxies.

### Role of High-End Computing (Why HPC Matters)

We have used 10s of millions of cpu-hours on NASA's Pleiades supercomputer to model both the large-scale distribution of galaxies in the cosmic web, and also regions including a central galaxy in its cosmic context. We have run the largest suite of galaxy simulations with path-breaking resolution and realistic astrophysical assumptions, and we have stored complete information at up to 2000 time steps. We use our Sunrise software to make realistic images and spectra of the simulated galaxies, including the effects of stellar evolution and dust scattering and absorption of starlight, for comparison with Hubble images.

### What's Next

We have recently greatly sped up our simulation code, which is allowing us to run more simulations all the way from the Big Bang to the present, with higher resolution and more realistic astrophysics. We are now simulating a wider range of galaxies, to understand the effects of early vs. late formation, and rapid vs. slow rotation, for example.

Presented at SC15 by Principal Investigator:

Joel Primack, University of California Santa Cruz, joel@ucsc.edu

Exhibit Theme: The Universe

## **Global Atmospheric Modeling on the Leading Edge of NASA Computing**

### **Overview:**

Global weather and climate models have evolved dramatically from their origins as basic mathematical models in the 1950s and 1960s. The growth and availability of more advanced computing over the latter part of the twentieth century lead to interactive atmosphere and land models and eventually fully coupled ocean/atmosphere Earth system models. The twentieth century has seen these global climate/weather models evolve into massive numerical missions including more components of the Earth system and representing more processes at much finer scales. Throughout this evolution of development, scientists have been willing explore the boundaries of computational capacity to push these models beyond their limitations. Often times scientists are willing to explore at a rate of just a single simulation day per day to see features never before seen in these models.

### **Project Details:**

At Goddard's NASA Center for Climate Simulation (NCCS) and the Global Modeling and Assimilation Office (GMAO) the development of the Goddard Earth Observing System atmospheric model (GEOS-5) over the last 10-years serves as a microcosm of this evolution. From spiraling storms in the tropics, to the fidelity of clouds over the North Pacific this global atmospheric model has evolved into an Earth system simulator depicting global weather and climate at resolutions never before explored on a global scale. GEOS-5 has gone from a collection of hundreds of thousands of grid cells covering the world in blocks more than 30 miles across, to hundreds of millions of grid cells just 1 mile wide. This evolution takes us from Hurricane Katrina in 2005 to Superstorm Sandy in 2012. We will explore stratocumulus clouds across the North Pacific and tornadoes in the Midwest United States. With branches along the way to explore fine particles across the globe, industrial and volcanic gases flowing from their sources, and the first global view of waves of carbon dioxide leaving their sources and engulfing the world.

Presented at SC15 by Principal Investigator:

William M Putman, NASA Goddard Space Flight Center, [william.m.putman@nasa.gov](mailto:william.m.putman@nasa.gov)

Exhibit Theme: Our Planet

## Predicting Damage from an Asteroid Strike on Earth

### Overview

The Spaceguard Survey has discovered approximately 10 000 Near Earth Asteroids larger than 30m, as of 2015. There are estimated to be over 1 million. Of those discovered 1600 are classified as hazardous. An asteroid larger than 30m is expected to strike the Earth on average every 100 years. NASA is investigating the consequences of an asteroid or comet strike on Earth to determine how large an asteroid should be dealt with by simple evacuation of an area, or whether it requires in-space mitigation such as trajectory alteration to miss the Earth or disruption to break it into small pieces. NASA will quantify how the mitigation strategy should depend on the asteroid properties such as size, density, speed, strength, and trajectory.

### Project Details

Prior to the asteroid strike in 2013 over the Russian city of Chelyabinsk, 20m diameter asteroids were not considered to be hazardous, as thankfully there is a great paucity of data on hazardous asteroid strikes on Earth. Although there are many meteor showers per year and many impact craters have been identified on Earth, the only well documented significant impact is the Chelyabinsk event. Consequently most asteroid damage models to date use simple equations describing the break-up of the asteroid in the atmosphere and analogies to nuclear explosions which were well studied in the 1950s. In lieu of sufficient observational data, computer simulations are being conducted to determine the asteroid fragmentation in the atmosphere or impact with the ground. Bursts in the air or ground emit pressure blast waves that propagate out and damage the surrounding area, and heat from the explosion can ignite fires.

### Results and Impact

Simulation of the Chelyabinsk event with a simple strength model of the rock, show the rock fractures when it hits the stratosphere. Rapid fragmentation and deceleration in the stratosphere produces energy deposition curves very similar to that observed from dash-cam videos of the event. Propagation of the resulting blast wave to the ground using a fluid dynamics code produces a ground damage very similar to that observed for broken windows. Improvements of the simulation to include thermal effects and application to different sizes and speeds will be used to fill out the range of possible impact scenarios that could occur. These results will be used to improve the analytical models and provide insight to mitigation recommendations.

### Role of High-End Computing (Why HPC Matters)

Simulation of an asteroid impact requires detailed modeling of the asteroid itself including its composition and how the rock mass behaves under the stresses associated with atmospheric entry and ground impact, and how the friction and heat erode and ablate the surface. To accurately simulate the behavior of the asteroid requires time steps of microseconds. The entry trajectory occurs over an altitude range of at least tens of kilometers, and damaging blast waves can extend hundreds of kilometers taking minutes to reach that distance. To capture the wide range of time and distance scales requires a supercomputer. A small 2D simulation just to capture the energy deposition rate can take as little as 100 000 grid cells and traverse the atmosphere in as little as 2000 processor-hours. Propagation of the blast wave to the ground takes 100 million grid cells and 10 000 processor-hours, at ~20 GFLOPS per processor.

### What's Next

The simulations are part of an asteroid threat assessment project to characterize asteroid properties, model the impacts, determine the sensitivity to parameters, and provide information of use to decision makers on future characterization and observation efforts, and mitigation strategies.

Presented at SC15 by:

Darrel Robertson, NASA Ames Research Center, [Darrel.K.Robertson@NASA.gov](mailto:Darrel.K.Robertson@NASA.gov)

Principal Investigator:

Donovan Mathias, NASA Ames Research Center, [Donovan.Mathias@NASA.gov](mailto:Donovan.Mathias@NASA.gov)

Exhibit Theme: Our Planet

## **IOPS Galore: Upgrading a Supercomputer's Metadata with Solid-State Drives**

### **Overview**

The NASA Center for Climate Simulation (NCCS) continually strives to provide faster, more advanced technology for its users. A part of the most recent wave of upgrades is a refresh of the metadata storage, which houses information such as inodes and directory tree structures. This refresh has increased the metadata capacity of the Discover supercomputer's storage system to 16 terabytes, significantly expanded user and administrator capabilities in metadata operations, and introduced new technology into the storage system.

### **Project Details**

Discover currently contains nearly 3,300 compute nodes and over 25 petabytes of usable storage. To support such a thriving operation, the NCCS physically separates data storage from metadata storage. Separating these storage types has many benefits, one of which is using different technologies that may not be feasible for the entire storage cluster.

The NCCS has taken advantage of this method by upgrading the previous spinning disk metadata storage to solid-state drives (SSDs). These drives are technologically superior to spinning disk for this purpose, with benefits including significantly higher capacity for input/output operations per second (IOPS) and avoiding common issues like "head thrashing" (and the slow file access that comes with it).

### **Results and Impact**

System administrators successfully transitioned all of Discover's metadata over to SSDs. They not only see benefits in administrative operations, but have multiple success stories from users as well. Improvements include metadata performance improvement of up to 10x and zero noticeable impact during rebuilds after failed disks.

### **Role of High-End Computing (Why HPC Matters)**

A high-performance computing environment gives scientists and engineers a place to generate and analyze large quantities of data within a reasonable amount of time. By providing newer and more advanced storage technologies, the NCCS enables these researchers to work with their data at a fraction of the time cost, bringing greater and more complex scientific goals into reality.

Presented at SC15 by:

Jordan Robertson, NASA Goddard Space Flight Center [jordan.robertson@nasa.gov](mailto:jordan.robertson@nasa.gov)

Principal Investigator:

Bruce Pfaff, NASA Goddard Space Flight Center, [bruce.pfaff@nasa.gov](mailto:bruce.pfaff@nasa.gov)

Exhibit Theme: Supercomputing

## Using CFD To Build The Space Launch System Booster Separation Database

### Overview

NASA's new Space Launch System (SLS) will be the first rocket since the Saturn V to carry astronauts beyond low-earth orbit, carrying 10% more payload than Saturn V and three times the Space Shuttle payload. The SLS configuration consists of a center core and two solid rocket boosters. The two boosters separation from the core as their fuel is exhausted two minutes from lift-off. There are 16 booster-separation motors (BSMs) whose short-duration burn provide the impulse to push the boosters away from the core. The aerodynamic forces play an important role in determining the flight of the boosters away from the core. Therefore an accurate booster-separation aerodynamic database is needed when verifying the design of the booster-separation system. This database is being built with a combination of computational fluid dynamics (CFD) simulations and a wind-tunnel test.

### Project Details

The goal of the current work is to build an aerodynamic database for the SLS booster-separation event. The approach is to use CFD simulations with the inviscid Cart3D flow solver to compute the aerodynamic forces on the two solid-rocket boosters and the core. Additional terms for the database were built using the results from a wind-tunnel test and with the high-fidelity Overflow CFD solver. The CFD simulations include simulating the aerodynamic effects of the 22 different plumes which are firing during the separation. These include the 16 BSM plumes, the plumes from the four core-stage main engines, and two plumes from the boosters.

The aerodynamic data is a function of eight independent variables. These include three translation variables and two rotation variables to define the booster position and orientation relative to the core. They include the free-stream angle of attack and angle of side-slip of the core, and the relative thrust of the booster-separation motors. In addition, separate sets of computations were run to simulate the effect when one of the core-stage engines failed.

### Results and Impact

Over 22,000 Cart3D simulations were to fill the parametric space during the booster separation event. In addition, 390 Overflow simulations were run to provide higher-fidelity data. Comparisons between the Cart3D results, the wind-tunnel data, and the Overflow data showed that the Cart3D code could accurately predict the booster and core aerodynamics to within an acceptable margin.

### Role of High-End Computing Resources

All of the CFD simulations were run on the pleiades supercomputer, primarily on the 24-core Haswell nodes. Each of the Cart3D simulations required approximately 200 core hours, and the Overflow runs required at least 20,000 core hours. A total of 12.7 million core hours were used to complete the computations. Most of these simulations were run over a two month period during which the project utilized 400 dedicated Haswell nodes (9600 cores). During this time the group used approximately 50 terabytes of short-term disk storage.

### What's Next

The new SLS booster-separation database will be used by the Guidance, Navigation, and Control group to verify the design of the SLS booster-separation system.

Presented at SC15 by Principal Investigator:

Stuart Rogers, NASA Ames Research Center, [stuart.e.rogers@nasa.gov](mailto:stuart.e.rogers@nasa.gov)

Exhibit Theme: Space Exploration

## Predictability of Hurricanes Katrina and Sandy: Scale or Detail Matters?

### Overview

*“NASA’s role as a research agency is to bring new types of observational capabilities and analytical tools to learn about the fundamental processes that drive hurricanes and work to help incorporate that data into forecasts.”*

Recent advances in supercomputing have made significant improvement in global climate modeling and the accompanying remarkable simulations now enable us to revisit the fundamental role of butterfly effect in the predictability, which is foundation of the chaos theory. Research activities have been conducted to improve our understanding of whether and how the scales and/or details of data, i.e., multiscale processes and butterfly effect, may impact dynamics of hurricanes and thus predictions of hurricanes.

### Project Details

We have developed a scalable, multiscale analysis package (MAP), which consists of the parallel ensemble empirical mode decomposition (PEEMD) and stability analysis tool (SAT), and integrated it into the Coupled Advanced multiscale Modeling and Visualization system (CAMVis) to improve extended-range tropical cyclone (TC) prediction and consequently TC climate projection by:

- Understanding the TC genesis processes, accompanying multiscale processes (both downscaling by large-scale events and upscaling by small-scale events), and their subsequent non-linear interactions
- Discovering the hidden predictive relationships among mesoscale TCs and large-scale tropical waves as well as small-scale convective processes
- Identifying the downscaling and upscaling processes in mathematical equations and numerical models that may form a nonlinear feedback loop to stabilize or destabilize solutions

### Results and Impact

By running the CAMVis and MAP on Pleiades supercomputer, we have made the following achievements:

- The CAMVis produced encouraging 5-day track and intensity predictions and insightful visualizations for Hurricane Katrina, showing its multiscale interactions with the approaching upper-level trough that may lead to the Katrina’s intensification prior to its landfall;
- The CAMVis produced remarkable 6~7-day track and genesis simulations of Hurricane Sandy; Multiscale analysis with the MAP/PEEMD and multiscale visualizations reveal the impact of tropical waves and upper-level trough on the formation and movement of Sandy
- Nonlinear analysis with the SAT in newly derived high-order Lorenz models suggest that negative nonlinear feedback can suppress chaotic responses and thus improve solution’s stability, suggesting properly decreasing grid spacing may reduce chaotic responses
- While the integrated system CAMVis produced remarkable TC simulations, with the MAP (PEEMD and SAT) revealed insightful multiscale processes that contribute to the TC’s predictability.

### Role of High-End Computing (Why HPC Matters)

Recent advances in supercomputing power has made it feasible to deploy the CAMVis and MAP. *The advanced, sophisticated, integrative system enables the research cycle of Modeling, Observation, Analysis, Synthesis and Theorizing (MOAST)*, illustrating (1) how research results from models and satellite are inter-compared (2) how multiscale analysis and synthesis can be used for proposing hypothesis of TC genesis and (3) how the combined impact of (1) and (2) can help us revisit the role of butterfly effect in predictability. The research cycle cannot be completed without the seamlessly integrative supercomputer facilities with large-scale computing, visualization and high-speed and high-volume storage systems.

### What’s Next

Multi-year simulations for TC genesis are being performed and analyzed to understand how downscaling and upscaling processes may contribute the predictability of TCs in numerical models.

Presented at SC15 by Principal Investigator:

Bo-wen Shen, San Diego State University, bshen@mail.sdsu.edu

Exhibit Theme: Our Planet

## **GRAIL Gravity Maps Enable Unprecedented Insights into Moon's Crust and Interior**

### **Overview**

Although the Gravity Recovery and Interior Laboratory's (GRAIL) twin Ebb and Flow spacecraft orbited the Moon at 3,600 miles per hour, their instruments made precise measurements of slight changes in distance between the spacecraft, caused by differences in the lunar gravity field underneath.

Using these measurements, GRAIL researchers are producing the highest-resolution satellite-based gravity maps of any planet or natural satellite. Because variations in lunar gravity are caused both by visible features, such as mountains and craters, and by masses hidden beneath the lunar surface, these gravity maps are divulging unprecedented detail about the Moon's crust and interior.

### **Project Details**

Researchers use the 7 million tracking observations between GRAIL's two spacecraft to calculate precise spacecraft orbits and geodetic characteristics of the Moon. Then the GRAIL team makes successive runs of gravity field models, using increasing degree and order of spherical harmonics, to estimate errors and refine parameters for subsequent runs. The recent GRGM900C gravity maps resolve 6-km spatial blocks at the Moon's equator, with even higher resolution at the poles.

### **Results and Impact**

High-resolution GRAIL gravity maps show very strong correlation with topography measured by the Lunar Orbiter Laser Altimeter instrument on NASA's Lunar Reconnaissance Orbiter, with differences between the two revealing subsurface details. Combining high-resolution GRAIL gravity maps with other techniques, researchers have discovered the following:

- The Moon's crust is less dense and more porous and fractured than previously thought.
- Gravity anomalies from GRAIL maps reveal that the outline of the dark volcanic Oceanus Procellarum on the near side is bounded by a rectangular formation of lava-flooded rift valleys not visible on the Moon's surface.
- While the Moon's near side experienced the same asteroid bombardment as the far side, it developed very different features because the near side crust was hotter and thinner at the time of the impacts.

### **Role of High-End Computing (Why HPC Matters)**

Using spherical harmonic functions of degree and order 900, GRAIL's models consume 4,080 processor-cores and 6 terabytes of memory on the NASA Center for Climate Simulation (NCCS) Discover supercomputer. Completing the iterations for one map takes several months. More than 4 million processor-hours were consumed by the GRAIL runs, culminating in the complete degree 900 lunar gravity map. Among other support, the NCCS has assisted by fine-tuning the team's software code, expediting job execution, and providing separate disk nodes optimized for maximum input/output speeds.

### **What's Next**

Analysis of current results indicates researchers have not exhausted the detail present in the observations, and the GRAIL team is preparing higher-resolution runs to fully extract the gravity information inherent in GRAIL tracking data.

Presented at SC15 by:

Ellen Salmon, NASA Goddard Space Flight Center, [ellen.m.salmon@nasa.gov](mailto:ellen.m.salmon@nasa.gov)

Principal Investigator:

Sander Goossens, University of Maryland, Baltimore County and NASA Goddard Space Flight Center,  
[sander.j.goossens@nasa.gov](mailto:sander.j.goossens@nasa.gov)

Frank Lemoine, NASA Goddard Space Flight Center, [frank.g.lemoine@nasa.gov](mailto:frank.g.lemoine@nasa.gov)

Exhibit Theme: The Universe

## **Designing a New Spacecraft for Deep Space Exploration**

### **Overview**

NASA is designing a new spacecraft that will enable astronauts to explore various destinations in our solar system. The Orion Multi-Purpose Crew Vehicle (MPCV) is a highly flexible and robust spacecraft that will usher in a new era of deep space exploration. Combined with the Space Launch System (SLS), the Orion spacecraft will transport astronauts to asteroids, Mars, or other celestial bodies and return them safely back to Earth. In designing the multipurpose vehicle, engineers needed to analyze a wide range of entry conditions and verify that the spacecraft will safely protect its crew from the harsh environment of an atmospheric entry. Hundreds of high-fidelity computer simulations were performed to assist engineers in the design of the Thermal Protection System (TPS) and other key components of the Orion spacecraft.

### **Project Details**

Computational Fluid Dynamics (CFD) simulations are computed to estimate the exchange of energy between the Orion capsule and its environment at different flight conditions (such as entry speed, altitude, and orientation of the vehicle). In addition, computer simulations are used to model situations where high heating may occur; for instance, the interaction of the reaction control system (RCS) jets with the incoming flow or enhanced heating downstream of a compression pad. Numerical simulations at both flight and wind tunnel conditions are computed and compared with experimental data. Spacecraft designers then used the computational results and test data to select and to size the TPS materials on the Orion capsule.

### **Results and Impact**

Comparisons are underway to verify the accuracy of CFD simulations with the Exploration Flight Test (EFT-1) measurements and wind tunnel data. Initial assessments indicate good agreement between the numerical results and the experimental data. The comparisons will help engineers to validate and refine their aerothermal models and to reduce margins in their designs. These improvements will aid engineers in optimizing the design of future spacecrafts for the Exploration Missions (EM-1 and EM-2).

### **Role of High-End Computing (Why HPC Matters)**

Typically, each computer simulation for the Orion capsule required 300-400 processors and 2-3 days of computing time on the Pleiades supercomputer cluster at NASA Ames Research Center. Thanks to the vast capabilities offered by the NASA Advanced Supercomputing (NAS) facility, hundreds of these calculations were completed in a month.

The ability to rapidly turnaround solutions is critical in keeping up with an ambitious schedule outlined for the Orion MPCV/SLS program. High-end computing also allowed engineers to run detailed simulations of RCS jets and compression pads on large grids (upwards of 400 million grid points). The high-speed networks also made it possible for transferring large datasets between NASA centers at Ames Research Center, Johnson Space Center, and Langley Research Center.

Presented at SC15 by:

Chun Tang, NASA Ames Research Center, Chun.Y.Tang@nasa.gov

Principal Investigator:

Joseph Olejniczak, NASA Ames Research Center, Joseph.Olejniczak@nasa.gov

Exhibit Theme: Space Exploration



## HECC – Moving to the Future, Today

### Overview

The High-End Computing Capability Project provides an agency-wide compute environment focused on enabling science and engineering discovery across all technical mission directorates. Meeting that challenge requires an environment that is reliable, stable and expandable. This year we were able to close to double the on-line Lustre storage capacity and upgrade the archive system to the next generation tape technology while taking the major compute system, Pleiades to 5.3 petaFLOPS.

Expansion in HECC is currently limited by power and cooling constraints. To meet agency requirements, NASA headquarters and HECC are evaluating facility expansion options. A study is being conducted to evaluate multiple approaches to expansion and a prototype is being fielded to evaluate performance of a Modular High-performance computing Center (MHC) over a year of operation.

### Project Details

HECC meets the competing High End Computing requirements of the agency through a customer-focused team of professionals providing world-class services to a broad array of technical users. HECC engineered a solution comprised of hardware and services to meet the demands of capacity, capability and time-critical computing simultaneously to our users by remembering that we are here to enable their goals.

This year, HECC made a large procurement growing the main computing system from 4.5 PF to 5.3 PF and almost doubling the Lustre storage from 20 PB to almost 39 PB. The archive system was also upgraded, moving to the LTO-6 tape technology.

Maintaining a healthy growth profile became challenging as we hit the power and cooling limits of our existing facility. To address these limitations, both NASA headquarters and the HECC team began a process to evaluate the best path forward for the agency. A study is being run by headquarters to look at the broad array of options and HECC procured an MHC to aide in the evaluation of a new approach of fielding high performance computers. The team remains driven to provide the best high-end computing environment to our user community through advances in hardware and improved service delivery.

### Results and Impact

Pleiades moved to its 6<sup>th</sup> generation of Intel Xeon architecture with the installation of the Haswell racks. The system has been able to move through the different technologies advancing in processor and network interfaces while delivering the equivalent of over 500,000,000 Westmere node hours. Overall, the HECC project has had a very successful year.

- Pleiades was ranked 11th in the June Top 500 list and 5th on the HPCG list, a benchmark that more closely models the types of work done by our user community.
- Expansions in the storage areas brought our overall Lustre raw capacity to over 35 petaBytes.
- Installed LTO-6 tape drives as an interim move to the soon to be released LTO-7 technology
- Began a study of facility expansion options
- Purchased a prototype MHC to evaluate electricity and water required to cool a functioning high-performance computer throughout the year.

### What's Next

HECC will continue to strive to meet the demands of the NASA user community. The facility expansion options have the potential to have a huge upside in the project's ability to meet these demands.

Presented at SC15 by Principal Investigator:

William Thigpen, NASA Ames Research Center, [william.w.thigpen@nasa.gov](mailto:william.w.thigpen@nasa.gov)

Exhibit Theme: Supercomputing

## Counting Trees and Shrubs on the South Side of the Sahara using Cloud Computing

### **Abstract**

As the combinations of remote sensing observations and model outputs have grown, scientists are increasingly burdened with both the necessity and complexity of large-scale data analysis. Scientists are increasingly applying traditional high performance computing (HPC) solutions to solve their "Big Data" problems. While this approach has the benefit of limiting data movement, HPC systems are optimized for large-scale simulations and not for data analysis.

To solve these issues and to alleviate some of the strain on the HPC environment, the NASA Center for Climate Simulation (NCCS) has created the Advanced Data Analytics Platform (ADAPT). ADAPT combines high performance computing and virtualization technologies to create an on-site private cloud specifically designed for large-scale data analytics. The NCCS has already deployed this system for use with a variety of Goddard research projects, including the Arctic Boreal Vulnerability Experiment (ABoVE). In order to extend this approach even further and to make this system as flexible as possible, the NCCS has coupled this private environment with a commercial cloud, namely Amazon Web Services (AWS). The result is a highly extensible and elastic hybrid cloud environment where users can start using systems and services at Goddard and burst into commercial clouds as needed.

In order to fully develop the hybrid cloud approach, the NCCS partnered with Dr. Compton Tucker (NASA Goddard) and Dr. Paul Morin (Univ. of Minnesota) on a project to estimate the tree and bush biomass over the entire arid and semi-arid zone on the south side of the Sahara. Drs. Tucker and Morin submitted a proposal to the Intel Head in the Clouds Challenge and were subsequently awarded a research grant for compute and storage capabilities in AWS. In order to estimate the amount of carbon stored and create a baseline for future research, approximately 100 terabytes of high resolution National Geospatial-Intelligence Agency (NGA) imagery must be analyzed. Since the ADAPT system cannot be dedicated to this project and must meet the requirements for several research teams, it will take approximately ten months of time to analyze all the data on NASA resources. Through the use of the hybrid cloud approach of combining local resources with AWS, the goal of this project is to finish the analysis in one month of time. This talk will motivate the need for hybrid cloud resources within NASA and highlight the information technology infrastructure being put in place that will dramatically speed up science.

## **Billion-Node Graph: NEX Data Production, Provenance and Analytics Pipeline**

### **Overview**

The NASA Earth Exchange (NEX) is a supercomputing platform for the Earth science community that provides a mechanism for scientific collaboration and knowledge sharing. NEX combines state-of-the-art supercomputing, Earth system modeling, workflow and knowledge management, and NASA remote sensing data feeds to deliver a complete work environment in which users can explore and analyze large datasets, run modeling codes, collaborate on new or existing projects, and quickly share results among the Earth science communities.

In order to support increasing number of projects aiming to develop new science data products at unprecedented scales, the NEX team has developed an HPC/Cloud hybrid data production system that represents both the production rules and the execution of these rules as a directed graph. The graph can be analyzed and used to reproduce scientific results as well as integrated with NEX Knowledge Graph for improved search results.

### **Project Details**

In order to support and scale up scientific research, the NEX team has developed system that supports definition and deployment of production workflows, analysis of the performance of individual workflow components, quality monitoring and scientific reproducibility. The system enables research teams to build complex non-linear production pipelines and is already deployed on number of projects ranging from application of machine learning to counting trees across the entire continent to developing consistent monthly and annual global higher-level satellite products at 30m spatial resolution – first of its kind at this resolution.

Through semantic integration with the NEX knowledge graph and data management system the system can answer additional queries about data quality, resource usage or related publications. The output of the science data pipeline is a data product and a graph that represents the entire provenance of each individual data granule. In many cases this graph can reach into 100's of millions of nodes.

### **Results and Impact**

The system makes it possible to execute projects that would not have been possible in the past by small science teams. Such as the ongoing:

- 37-step production pipeline for Landsat and NASA Global Imagery and Browse Service – 7PB of data/670,000 compute hours
- NASA NEX-DCP-30 and NEX-GDDP downscaled climate projections pipeline – 40TB of data/320,000 compute hours
- Estimating biomass by counting trees across United States by using machine learning and computer vision – 40TB of data/100,000 compute hours

### **Role of High-End Computing (Why HPC Matters)**

The Pleiades supercomputing architecture combined with the massive data store and high-speed network enables NEX to engage large scientific communities and provide them with capabilities to execute modeling and data analysis on a grand scale, which was not previously achievable by most scientists. During the last year NEX projects have used over 860,000 hours on the Pleiades supercomputer and processed over 1PB of data.

### **What's Next**

We plan to both engage users and make it substantially easier to build Big Data production pipelines in the HPC environment. We also continue to scale up our engagement with machine learning and data mining community, provide support for long-term data records with Landsat, prototyping production system for Sentinel-2 and extending data pipeline to NASA Global Imagery and Browse Service (GIBS).

Presented at SC15 by:

Petr Votava, NASA Ames Research Center, Petr.Votava@nasa.gov

Principal Investigator:

Andrew Michaelis, NASA Ames Research Center, Andrew.R.Michaelis@nasa.gov

Exhibit Theme: Our Planet

## **Climate Modeling 101: From NASA Observational Data to Climate Projections**

### **Overview**

Computer models that simulate the climate are an integral part of providing climate information, in particular for future changes in the climate. NASA, along with other National and International Government Agencies, Universities and industry participants are developing and testing Climate Models.

A climate model is a system of models (or sub models) that are integrated to simulate the Earth system. The sub models include atmospheric, oceanic, land, surface and sea ice, chemistry, hydrology and others. Each of these components represent different scientific disciplines may have different computational algorithms with individual characteristic time and length scales.

### **Use of NASA Satellite data:**

NASA invests significant resources to launch Earth Observing satellites to gather data. There are currently over two-dozen missions and NASA makes this data freely available to users around the world (figure 1). This data is used by modelers to understand physical phenomena and develop mathematical models of individual components, test the accuracy of the component models and test the accuracy of the integrated climate model.

### **Climate Models**

A climate model can generally be described as a set of time marching, nonlinear, partial differential equations. These equations are then cast into difference equations that can be solved numerically on a computer. The difference equations are solved at the nodes on a “numerical mesh”. There are three meshes, one for the atmospheric, one for the ocean and one for land surface. At the boundaries between the meshes, boundary conditions must be passed between the various sub-models.

Climate model runs are very large and in order to solve them, advanced computational techniques are required in order to use supercomputers. This includes techniques such as Message Passing Interface (MPI) to use thousands of cores.

### **Role of Supercomputing resources**

Climate models require supercomputers to run. A typical climate science project requires multi component ensembles, each using thousands of cores running for weeks at a time resulting in millions of core hours. Such a project can create hundreds of terabytes of data.

Multinational experiments are be conducted by the International Panel for Climate Change (IPCC). In this case, modelers at many different Intutions around the world run a set of numerical experiments, each using their own climate model. The most recent experiment CMIP-5 in 2013 created over 50 PBytes of data for the international community.

### **Results**

Results from Global Modeling and Assimilation Office (GMAO), Goddard institute for Space Studies (GISS) for weather and climate models will be presented.

Presented at SC15 by Principal Investigator:

W. Phil Webster, Goddard Space Flight Center, Phil.Webster@NASA.gov

Exhibit Theme: Our Planet

## **Simulation of the SLS Ignition Over-Pressure/Sound Suppression Water System**

### **Overview**

The Space Launch System (SLS) Launch vehicle consists of 4 RS-25 Liquid Rocket Engines and two five-segment Solid Rocket Boosters. Upon Ignition, the Solid Rocket Boosters create a significant overpressure event, the Ignition Over-Pressure (IOP) event. The IOP event experienced in the first Space Shuttle Flight, STS-1 in 1981, was significant enough to damage the Space Shuttle Orbiter and motivate significant design changes to the Space Shuttle Launch Pad. The RS-25 liquid rocket engines, along with the Solid Rocket Boosters, also create a very significant acoustic environment that, if not mitigated, can harm the SLS Launch Vehicle. The purpose of the Ignition Over-Pressure/Sound Suppression water system is to reduce the amplitude of both the IOP event and the acoustic environments caused by both the RS-25 engines and the Solid Rocket Boosters.

Mitigation of the IOP and Acoustic environments by the IOP/SS water system is dependent upon proper design of the IOP/SS system. One of several important aspects of the design is the correct placement of the water such that the water mitigates the environment while not interfering with the nominal operation of other subsystems in the vicinity such as the RS-25 engines, the Solid Rocket Booster and Hydrogen Burn-off Igniter (HBOI) Systems. The Volume of Fluid (VOF) model recently developed within the Loci-STREAM Computational Fluid Dynamics (CFD) program provides the ability to simulate the placement of water by the IOP/SS system.

### **Project Details**

A detailed CFD model of the SLS Launch Pad, including the IOP/SS water system and the lower portions of the SLS Launch Vehicle, was constructed. This model consisted of 82 million volume cells. The Loci-STREAM-VOF CFD program was used to perform an unsteady simulation of the IOP/SS water system. This first of its kind CFD simulation was used to determine the placement of the water with respect to the systems intended function and with respect to operation of other subsystems in the vicinity of the IOP/SS water system.

### **Results and Impact**

The CFD simulations of the initial design(s) identified water projection patterns which adversely interacted with both the Solid Rocket Boosters and the HBOI subsystems. Based on the adverse interactions with the Solid Rocket Boosters identified by the unsteady CFD simulations, the design of the IOP/SS system was significantly altered to avoid the problematic water projection while still performing its intended environments mitigation function. The CFD simulation also determined the IOP/SS system was creating an air entrainment effect, or waterfall effect, that interfered with the intended operation of the HBOI subsystem. In this case, an HBOI system design change was devised that accommodated the IOP/SS effect. As a result of this work, the IOP/SS and HBOI systems are expected to operate as intended as opposed to detecting these problems much later in full-scale testing activities, thus avoiding expensive redesigns and schedule consequences.

### **Role of High-End Computing (Why HPC Matters)**

The Loci-STREAM-VOF program was developed with the explicit purpose of enabling the use of high performance computing for the Volume of Fluid algorithm. The state of the art in parallel computing using the VOF algorithm for unstructured meshes was limited to approximately 32 cpu cores. The Loci-STREAM-VOF program was used to simulate the SLS IOP/SS system using over 4,000 concurrent cpu cores. An assessment of this kind would have not been possible without the NASA HPC hardware and the software written to exploit it.

### **What's Next**

The current CFD simulation simulates the placement of water by the IOP/SS system before the ignition of the liquid engines and solid rocket boosters. The next steps include the combination of the VOF water simulation with the simulation of the Solid Rocket Motor Ignition Over-Pressure event. The purpose of the simulation is to quantitatively predict the mitigation of the Ignition Over-Pressure event.

Presented at SC15 by Principal Investigator:

Jeff West, NASA Marshall Space Flight Center, jeffrey.s.west@nasa.gov

Exhibit Theme: Space Exploration

## **Analyzing the Space Launch System Debris Environment**

### **Overview**

Following the Space Shuttle Columbia tragedy, NASA has made a concerted effort to identify all sources of debris, such as ice and foam from the propellant tanks, and quantify the impact energies this debris may impart to the launch vehicle, which is called the debris environment. This process is now being applied during the design phase of NASA's Space Launch System (SLS). Once the debris and impacts are fully characterized, the SLS vehicle designers can assess whether their hardware is capable of surviving this environment, or try to eliminate the debris sources. An overly conservative analysis of the debris environment may drive expensive design changes which can ultimately reduce the launch capacity of the SLS vehicle. Therefore, NASA uses high-fidelity computational fluid dynamics (CFD) simulations to understand the fluid flow features that may act as transport mechanisms for debris. At liftoff, these features include the solid rocket booster ignition overpressure, rocket plume entrainment and impingement, and ambient winds. The CFD simulation data is then coupled with models of typical debris to predict the resulting debris trajectories, and identify the affected components of the vehicle.

### **Project Details**

Analysis of the SLS debris environment required a two-pronged approach to address the various fluid flow aspects at liftoff. First, a series of static vehicle CFD simulations was conducted to vary the wind speed and direction, and the on/off state of the rocket plumes. Depending on the debris type and source location, different wind speed/direction combinations may result in maximum impact conditions at different parts of the vehicle. Additionally, the entrainment from the liquid engine plumes and/or the solid rocket plumes can significantly alter the course of some types of debris near the aft end of the vehicle.

Second, a dynamic CFD simulation of the SLS vehicle moving on a launch trajectory was used to investigate plume-driven flow features during the first few seconds of flight. Plume-driven flow (including the solid rocket booster ignition overpressure, plume impingement and recirculation) is especially important because it has the potential to accelerate debris near the aft end of the vehicle and the launch platform deck to much greater speeds than wind or gravity alone. These flow features are truly transient as the vehicle lifts off the launch pad, and cannot be captured with static CFD simulations.

### **Results and Impact**

In all, 25 static simulations and the dynamic launch simulation were employed to analyze the liftoff flow field and the resulting debris environment for SLS. The use of high-fidelity models and simulations reduced the uncertainty and conservatism inherent in earlier debris environment analyses which used simpler engineering models. For example, the SLS launch simulation was interrogated for high-speed upward flow resulting from plume impingement and recirculation on the launch deck, which was a major factor in previous debris analyses, but no significant upward flow was identified. The delivery of the updated debris environment was important to the completion of SLS Critical Design Review earlier this year.

### **Role of High-End Computing (Why HPC Matters)**

NASA supercomputing resources are instrumental to completing this type of analysis. The computational model of the SLS vehicle and launch pad requires nearly 250 million mesh cells to adequately represent the effects of the rocket plumes and the ambient winds. Each of the static simulations in this work consumed approximately 75,000 CPU-hours. The dynamic simulation of the SLS launch used nearly 4 million CPU-hours, which would have been an intractable problem even 10 years ago during the Space Shuttle Program. Now, we can leverage NASA's HPC resources to field these types of high-fidelity simulations during the SLS design process.

Presented at SC15 by Principal Investigator:

Brandon Williams, NASA Marshall Space Flight Center, [brandon.williams@nasa.gov](mailto:brandon.williams@nasa.gov)

Exhibit Theme: Space Exploration

## **Study of Baffled Propellant Tank Slosh using Multiphase CFD**

### **Overview**

When liquid propellant slosh frequency coincides with either the fundamental elastic body bending frequency or the dynamic control frequency of the vehicle at some time during the powered phase of the flight, the slosh forces could interact with the structure or control system. This could cause a failure of structural components within the vehicle or excessive deviation from the planned flight path. It is important therefore to consider means to provide adequate damping of the liquid motion and slosh forces and to develop methods to account for damping in the analysis of vehicle performance. For axisymmetric tanks, a common method of increasing damping is to attach a series of ring baffles to the tank wall. The baffles provide a substantial degree of damping when the free surface is near one of the baffles. The spacing between the baffles is chosen so that the damping exceeds the requirements regardless of fill level. Fuel-slosh damping by ring baffles in cylindrical tanks has been investigated extensively in the 1960s both experimentally and theoretically. A survey of damping measurements shows many apparent discrepancies. The objective of this effort is to improve our understanding of the physics behind slosh damping and to validate damping extracted from CFD against experimental data using the same parameters in a baffled tank.

### **Project Details**

Accurately predicting the slosh damping value of a baffled tank is very challenging for any Computational Fluid Dynamics (CFD) tool. One must resolve thin boundary layers near the wall and subscale vortices near the baffles. In this study, a Volume-Of-Fluid (VOF) module was linked to a production CFD solver at NASA/MSFC, Loci-STREAM. This new program enables a fast turn-around time and a lower numerical damping by using high mesh density enabled by the massively parallel processing capability of NASA's Pleiades supercomputing cluster. The CFD program was first validated against experimental data for slosh damping inside a smooth-wall tank. The slosh damping values for a baffled tank were then computed and compared to experiment.

### **Results and Impact**

The experimental data correlated into the industry standard for the smooth wall were used as the baseline validation. It was demonstrated that with proper grid resolution, CFD can indeed accurately predict low damping values from smooth walls for different tank sizes. The damping due to ring baffles at different depths from the free surface was then simulated, and very good agreement with experimental data was observed. This effort clearly illustrated the soundness of the CFD approach in modeling the detailed fluid dynamics of tank sloshing and the excellent accuracy in extracting slosh mechanical properties, especially sloshing damping.

Application of the CFD modeling tool has great potential in the design of propellant tanks with multiple baffles. In general, the practice to estimate the total damping for multiple baffles is by a linear superposition of the estimated damping contribution from each individual baffle. However, experimental data involving multiple baffles are limited. CFD simulation could be used to determine the validity and applicability of such rules. Finally, CFD simulation could also be used to help develop a relation between slosh frequencies and baffle parameters. Using baffles to shift frequency could be an alternative to tank compartmentalization.

### **Role of High-End Computing (Why HPC Matters)**

Capturing thin boundary profile and subscale vortices of the slosh wave requires very fine grid resolution that leads to huge memory requirements and intensive computation. Adding the complexity of the solution of two different fluids with a high-density ratio (1000 for water to air), the resulting equations become very stiff. The present simulations could not have been possible without the use of a supercomputer. All the simulations were performed on NASA's Pleiades supercomputer cluster using 500 to 1000 processors. NASA's High-Performance Computing (HPC) Pleiades cluster allows us to solve complex multi-phase flow problems through a high bandwidth, low latency network, providing highly parallel capabilities. The computational runs took 1 to 3 days to complete.

Presented at SC15 by Principal Investigator:

H. Q. Yang, NASA Marshall Space Flight Center, hong.q.yang@nasa.gov

Exhibit Theme: Space Exploration

## **Interactional Aerodynamics of Vertical-Lift Multi-Rotor Drones**

### **Overview**

Recent interest in technologies that will decrease emissions and fuel burn has led to research in aircraft employing hybrid or full electric propulsion systems. Multi-rotors are frequently used for small electric rotorcraft partly due to mechanical simplicity. The use of multiple rotors provides redundancy as well as cost-efficiency. However, unlike most helicopters with single main rotor, the interactions between multiple rotors and a fuselage emerge as an important factor to consider in design.

### **Project Details**

High-fidelity computational simulation tools are required to investigate the physics of multi-rotor interactions that are not well understood at this time. For applications to civil operations, it is necessary to assess aerodynamic performance and noise levels of multi-rotors. Present study covers not only small-scale drones but also large-scale multi-rotors for heavy-lifting missions.

### **Results and Impact**

- Computational simulations are performed using thousands of cores on the Pleiades supercomputer to analyze flow interactions between multiple rotors, wings, and a fuselage.
- New conceptual designs of large drones including a dual-winged quadcopter and a tandem coaxial rotorcraft are investigated.
- Various hybrid turbulence models including detached eddy simulation are assessed for interactional aerodynamics.

### **Role of High-End Computing (Why HPC Matters)**

Access to the Pleiades supercomputer systems has been essential for achieving the project goal. NAS visualization has been critically important to understand complex, unsteady, turbulent flow interactions. Tens of millions of CPU hours have been used.

Presented at SC15 by Principal Investigator:  
Seokkwan Yoon, NASA Ames Research Center, s.yoon@nasa.gov

Exhibit Theme: Aeronautics